arXiv:1401.4130v1 [cs.LO] 16 Jan 2014

Analysis of Probabilistic Basic Parallel Processes

Rémi Bonnet[1], Stefan Kiefer[1⋆], and Anthony W. Lin[1,2]

[1] University of Oxford, UK
[2] Academia Sinica, Taiwan

Abstract. Basic Parallel Processes (BPPs) are a well-known subclass of Petri Nets. They are the simplest common model of concurrent programs that allows unbounded spawning of processes. In the probabilistic version of BPPs, every process generates other processes according to a probability distribution. We study the decidability and complexity of fundamental qualitative problems over probabilistic BPPs — in particular reachability with probability 1 of different classes of target sets (e.g. upward-closed sets). Our results concern both the Markov-chain model, where processes are scheduled randomly, and the MDP model, where processes are picked by a scheduler.

1 Introduction

We study probabilistic basic parallel processes (pBPP), which is a stochastic model for concurrent systems with unbounded process spawning. Processes can be of different types, and each type has a fixed probability distribution for generating new sub-processes. A pBPP can be described using a notation similar to that of stochastic context-free grammars. For instance,

$$X \xhookrightarrow{0.2} XX \qquad X \xhookrightarrow{0.3} XY \qquad X \xhookrightarrow{0.5} \varepsilon \qquad\qquad Y \xhookrightarrow{0.7} X \qquad Y \xhookrightarrow{0.3} Y$$

describes a system with two types of processes. Processes of type X can generate two processes of type X, one process of each type, or zero processes with probabilities 0.2, 0.3, and 0.5, respectively. Processes of type Y can generate one process, of type X or Y, with probability 0.7 and 0.3. The order of processes on the right-hand side of each rule is not important. Readers familiar with process algebra will identify this notation as a probabilistic version of Basic Parallel Processes (BPPs), which is widely studied in automated verification, see e.g. [7, 11, 6, 13, 12, 9],

A *configuration* of a pBPP indicates, for each type X, how many processes of type X are present. Writing Γ for the finite set of types, a configuration is thus an element of $\mathbb{N}^\Gamma$. In a configuration $\alpha \in \mathbb{N}^\Gamma$ with $\alpha(X) \geq 1$ an X-process may be scheduled. Whenever a process of type X is scheduled, a rule with X on the left-hand side is picked randomly according to the probabilities of the rules, and then an X-process is replaced by processes as on the right-hand side.

⋆ Stefan Kiefer is supported by a Royal Society University Research Fellowship.

In the example above, if an X-process is scheduled, then with probability 0.3 it is replaced by a new X-process and by a new Y-process. This leads to a new configuration, α', with $\alpha'(X) = \alpha(X)$ and $\alpha'(Y) = \alpha(Y) + 1$.

Which type is scheduled in a configuration $\alpha \in \mathbb{N}^\Gamma$ depends on the model under consideration. One possibility is that the type to be scheduled is selected randomly among those types X with $\alpha(X) \geq 1$. In this way, a pBPP induces an (infinite-state) Markov chain. We consider two versions of this Markov chain: in one version the type to be scheduled is picked using a uniform distribution on those types with at least one waiting process; in the other version the type is picked using a uniform distribution on the waiting processes. For instance, in configuration α with $\alpha(X) = 1$ and $\alpha(Y) = 2$, according to the "type" version, the probability of scheduling X is $1/2$, whereas in the "process" version, the probability is $1/3$. Both models seem to make equal sense, so we consider them both in this paper. As it turns out their difference is unimportant for our results.

In many contexts (e.g. probabilistic distributed protocols — see [15, 14]), it is more natural that this scheduling decision is not taken randomly, but by a *scheduler*. Then the pBPP induces a Markov decision process (MDP), where a scheduler picks a type X to be scheduled, but the rule with X on the left-hand side is selected probabilistically according to the probabilities on the rules.

In this paper we provide decidability results concerning *coverability* with probability 1, or "almost-sure" coverability, which is a fundamental qualitative property of pBPPs. We say a configuration $\beta \in \mathbb{N}^\Gamma$ *covers* a configuration $\phi \in \mathbb{N}^\Gamma$ if $\beta \geq \phi$ holds, where $\geq$ is meant componentwise. For instance, ϕ may model a configuration with one producer and one consumer; then $\beta \geq \phi$ means that a transaction between a producer and a consumer can take place. Another example is a critical section that can be entered only when a lock is obtained. Given a pBPP, an initial configuration α, and target configurations $\phi_1, \ldots, \phi_k$, the *coverability problem* asks whether with probability 1 it is the case that starting from α a configuration β is reached that covers some ϕ_i. One can equivalently view the problem as almost-sure reachability of an upward-closed set.

In Section 3 we show using a Karp-Miller-style construction that the coverability problem for pBPP Markov chains is decidable. We provide a nonelementary lower complexity bound. In Section 4 we consider the coverability problem for MDPs. There the problem appears in two flavours, depending on whether the scheduler is "angelic" or "demonic". In the angelic case we ask whether there *exists* a scheduler so that a target is almost-surely covered. We show that this problem is decidable, and if such a scheduler does exist one can synthesize one. In the demonic case we ask whether a target is almost-surely covered, no matter what the scheduler (an operating system, for instance) does. For the question to make sense we need to exclude unfair schedulers, i.e., those that never schedule a waiting process. Using a robust fairness notion (k-fairness), which does not depend on the exact probabilities in the rules, we show that the demonic problem is also decidable. In Section 5 we show for the Markov chain and for both versions of the MDP problem that the coverability problem becomes P-time solvable, if the target configurations ϕ_i consist of only one process each (i.e., are

unit vectors). Such target configurations naturally arise in concurrent systems (e.g. freedom from deadlock: whether *at least one* process eventually goes into a critical section). Finally, in Section 6 we show that the almost-sure reachability problem for semilinear sets, which generalizes the coverability problem, is undecidable for pBPP Markov chains and MDPs. Some missing proofs can be found in the appendix.

Related work. (Probabilistic) BPPs can be viewed as (stochastic) Petri nets where each transition has exactly one input place. Stochastic Petri nets, in turn, are equivalent to probabilistic vector addition systems with states (pVASSs), whose reachability and coverability problems were studied in [1]. This work is close to ours; in fact, we build on fundamental results of [1]. Whereas we show that coverability for the Markov chain induced by a pBPP is decidable, it is shown in [1] that the problem is undecidable for general pVASSs. In [1] it is further shown for general pVASSs that coverability becomes decidable if the target sets are "Q-states". If we apply the same restriction on the target sets, coverability becomes polynomial-time decidable for pBPPs, see Section 5. MDP problems are not discussed in [1].

The MDP version of pBPPs was studied before under the name *task systems* [2]. There, the scheduler aims at a "space-efficient" scheduling, which is one where the maximal number of processes is minimised. Goals and techniques of this paper are very different from ours.

Certain classes of non-probabilistic 2-player games on Petri nets were studied in [16]. Our MDP problems can be viewed as games between two players, Scheduler and Probability. One of our proofs (the proof of Theorem 11) is inspired by proofs in [16].

The notion of k-fairness that we consider in this paper is not new. Similar notions have appeared in the literature of concurrent systems under the name of "bounded fairness" (e.g. see [5] and its citations).

2 Preliminaries

We write $\mathbb{N} = \{0, 1, 2, \ldots\}$. For a countable set X we write $\mathit{dist}(X)$ for the set of *probability distributions* over X; i.e., $\mathit{dist}(X)$ consists of those functions $f : X \to [0, 1]$ such that $\sum_{x \in X} f(x) = 1$.

Markov Chains. A *Markov chain* is a pair $\mathcal{M} = (Q, \delta)$, where Q is a countable (finite or infinite) set of states, and $\delta : Q \to \mathit{dist}(Q)$ is a probabilistic transition function that maps a state to a probability distribution over the successor states. Given a Markov chain we also write $s \xrightarrow{p} t$ or $s \to t$ to indicate that $p = \delta(s)(t) > 0$. A *run* is an infinite sequence $s_0 s_1 \cdots \in Q^\omega$ with $s_i \to s_{i+1}$ for $i \in \mathbb{N}$. We write $\mathit{Run}(s_0 \cdots s_k)$ for the set of runs that start with $s_0 \cdots s_k$. To every initial state $s_0 \in S$ we associate the probability space $(\mathit{Run}(s_0), \mathcal{F}, \mathcal{P})$ where $\mathcal{F}$ is the σ-field generated by all basic cylinders $\mathit{Run}(s_0 \cdots s_k)$ with $s_0 \cdots s_k \in Q^*$, and $\mathcal{P} : \mathcal{F} \to [0, 1]$ is the unique probability measure such that $\mathcal{P}(\mathit{Run}(s_0 \cdots s_k)) =$

$\prod_{i=1}^{k} \delta(s_{i-1})(s_i)$. For a state $s_0 \in Q$ and a set $F \subseteq Q$, we write $s_0 \models \Diamond F$ for the event that a run started in s_0 hits F. Formally, $s_0 \models \Diamond F$ can be seen as the set of runs $s_0 s_1 \cdots$ such that there is $i \geq 0$ with $s_i \in F$. Clearly we have $\mathcal{P}(s_0 \models \Diamond F) > 0$ if and only if in $\mathcal{M}$ there is a path from s_0 to a state in F. Similarly, for $Q_1, Q_2 \subseteq Q$ we write $s_0 \models Q_1 \mathsf{U} Q_2$ to denote the set of runs $s_0 s_1 \cdots$ such that there is $j \geq 0$ with $s_j \in Q_2$ and $s_i \in Q_1$ for all $i < j$. We have $\mathcal{P}(s_0 \models Q_1 \mathsf{U} Q_2) > 0$ if and only if in $\mathcal{M}$ there is a path from s_0 to a state in Q_2 using only states in Q_1. A Markov chain is *globally coarse* with respect to a set $F \subseteq Q$ of configurations, if there exists $c > 0$ such that for all $s_0 \in Q$ we have that $\mathcal{P}(s_0 \models \Diamond F) > 0$ implies $\mathcal{P}(s_0 \models \Diamond F) \geq c$.

Markov Decision Processes. A *Markov decision process (MDP)* is a tuple $\mathcal{D} = (Q, A, \mathit{En}, \delta)$, where Q is a countable set of states, A is a finite set of actions, $\mathit{En} : Q \to 2^A \setminus \emptyset$ is an action enabledness function that assigns to each state s the set $\mathit{En}(s)$ of actions enabled in s, and $\delta : S \times A \to \mathit{dist}(S)$ is a probabilistic transition function that maps a state s and an action $a \in \mathit{En}(s)$ enabled in s to a probability distribution over the successor states. A *run* is an infinite alternating sequence of states and actions $s_0 a_1 s_1 a_2 \cdots$ such that for all $i \geq 1$ we have $a_i \in \mathit{En}(s_{i-1})$ and $\delta(s_{i-1}, a_i)(s_i) > 0$. For a finite word $w = s_0 a_1 \cdots s_{k-1} a_k s_k \in Q(AQ)^*$ we write $\mathit{last}(w) = s_k$. A *scheduler* for $\mathcal{D}$ is a function $\sigma : Q(AQ)^* \to \mathit{dist}(A)$ that maps a run prefix $w \in Q(AQ)^*$, representing the history of a play, to a probability distribution over the actions enabled in $\mathit{last}(w)$. We write $\mathit{Run}(w)$ for the set of runs that start with $w \in Q(AQ)^*$. To an initial state $s_0 \in S$ and a scheduler σ we associate the probability space $(\mathit{Run}(s_0), \mathcal{F}, \mathcal{P}_\sigma)$, where $\mathcal{F}$ is the σ-field generated by all basic cylinders $\mathit{Run}(w)$ with $w \in \{s_0\}(AQ)^*$, and $\mathcal{P}_\sigma : \mathcal{F} \to [0, 1]$ is the unique probability measure such that $\mathcal{P}(\mathit{Run}(s_0)) = 1$, and $\mathcal{P}(\mathit{Run}(was)) = \mathcal{P}(\mathit{Run}(w)) \cdot \sigma(w)(a) \cdot \delta(\mathit{last}(w), a)(s)$ for all $w \in \{s_0\}(AQ)^*$ and all $a \in A$ and all $s \in Q$. A scheduler σ is called *deterministic* if for all $w \in Q(AQ)^*$ there is $a \in A$ with $\sigma(w)(a) = 1$. A scheduler σ is called *memoryless* if for all $w, w' \in Q(AQ)^*$ with $\mathit{last}(w) = \mathit{last}(w')$ we have $\sigma(w) = \sigma(w')$. When specifying events, i.e., measurable subsets of $\mathit{Run}(s_0)$, the actions are often irrelevant. Therefore, when we speak of runs $s_0 s_1 \cdots$ we mean the runs $s_0 a_1 s_1 a_2 \cdots$ for arbitrary $a_1, a_2, \ldots \in A$. E.g., in this understanding we view $s_0 \models \Diamond F$ with $s_0 \in Q$ and $F \subseteq Q$ as an event.

Probabilistic BPPs and their configurations. A *probabilistic BPP (pBPP)* is a tuple $\mathcal{S} = (\Gamma, \hookrightarrow, \mathit{Prob})$, where Γ is a finite set of *types*, $\hookrightarrow \subseteq \Gamma \times \mathbb{N}^\Gamma$ is a finite set of *rules* such that for every $X \in \Gamma$ there is at least one rule of the form $X \hookrightarrow \alpha$, and Prob is a function that to every rule $X \hookrightarrow \alpha$ assigns its probability $\mathit{Prob}(X \hookrightarrow \alpha) \in (0, 1] \cap \mathbb{Q}$ so that for all $X \in \Gamma$ we have $\sum_{X \hookrightarrow \alpha} \mathit{Prob}(X \hookrightarrow \alpha) = 1$. We write $X \overset{p}{\hookrightarrow} \alpha$ to denote that $\mathit{Prob}(X \hookrightarrow \alpha) = p$. A *configuration* of $\mathcal{S}$ is an element of $\mathbb{N}^\Gamma$. We write $\alpha_1 + \alpha_2$ and $\alpha_1 - \alpha_2$ for componentwise addition and subtraction of two configurations α_1, α_2. When there is no confusion, we may identify words $u \in \Gamma^*$ with the configuration $\alpha \in \mathbb{N}^\Gamma$ such that for all $X \in \Gamma$ we have that $\alpha(X) \in \mathbb{N}$ is the number of occurrences of X in u. For instance, we write XXY or XYX for the configuration α with $\alpha(X) = 2$ and $\alpha(Y) = 1$ and

$\alpha(Z) = 0$ for $Z \in \Gamma \setminus \{X, Y\}$. In particular, we may write ε for α with $\alpha(X) = 0$ for all $X \in \Gamma$. For configurations α, β we write $\alpha \leq \beta$ if $\alpha(X) \leq \beta(X)$ holds for all $X \in \Gamma$; we write $\alpha < \beta$ if $\alpha \leq \beta$ but $\alpha \neq \beta$. For a configuration α we define the *number of types* $|\alpha|_{type} = |\{X \in \Gamma \mid \alpha(X) \geq 1\}|$ and the *number of processes* $|\alpha|_{proc} = \sum_{X \in \Gamma} \alpha(X)$. Observe that we have $|\alpha|_{type} \leq |\alpha|_{proc}$. A set $F \subseteq \mathbb{N}^\Gamma$ of configurations is called *upward-closed* (*downward-closed*, respectively) if for all $\alpha \in F$ we have that $\alpha \leq \beta$ implies $\beta \in F$ ($\alpha \geq \beta$ implies $\beta \in F$, respectively). For $\alpha \in \mathbb{N}^\Gamma$ we define $\alpha{\uparrow} := \{\beta \in \mathbb{N}^\Gamma \mid \beta \geq \alpha\}$. For $F \subseteq \mathbb{N}^\Gamma$ and $\alpha \in F$ we say that α is a *minimal element* of F, if there is no $\beta \in F$ with $\beta < \alpha$. It follows from Dickson's lemma that every upward-closed set has finitely many minimal elements; i.e., F is upward-closed if and only if $F = \phi_1{\uparrow} \cup \ldots \cup \phi_n{\uparrow}$ holds for some $n \in \mathbb{N}$ and $\phi_1, \ldots, \phi_n \in \mathbb{N}^\Gamma$.

Markov Chains induced by a pBPP. To a pBPP $\mathcal{S} = (\Gamma, \hookrightarrow, Prob)$ we associate the Markov chains $\mathcal{M}_{type}(\mathcal{S}) = (\mathbb{N}^\Gamma, \delta_{type})$ and $\mathcal{M}_{proc}(\mathcal{S}) = (\mathbb{N}^\Gamma, \delta_{proc})$ with $\delta_{type}(\varepsilon, \varepsilon) = \delta_{proc}(\varepsilon, \varepsilon) = 1$ and for $\alpha \neq \varepsilon$

$$\delta_{type}(\alpha, \gamma) = \sum_{\substack{X \overset{p}{\hookrightarrow} \beta \text{ s.t. } \alpha(X) \geq 1 \\ \text{and } \gamma = \alpha - X + \beta}} \frac{p}{|\alpha|_{type}} \quad \text{and} \quad \delta_{proc}(\alpha, \gamma) = \sum_{\substack{X \overset{p}{\hookrightarrow} \beta \text{ s.t.} \\ \gamma = \alpha - X + \beta}} \frac{\alpha(X) \cdot p}{|\alpha|_{proc}}.$$

In words, the new configuration γ is obtained from α by replacing an X-process with a configuration randomly sampled according to the rules $X \overset{p}{\hookrightarrow} \beta$. In $\mathcal{M}_{type}(\mathcal{S})$ the selection of X is based on the number of types in α, whereas in $\mathcal{M}_{proc}(\mathcal{S})$ it is based on the number of processes in α. We have $\delta_{type}(\alpha, \gamma) = 0$ iff $\delta_{proc}(\alpha, \gamma) = 0$. We write $\mathcal{P}_{type}$ and $\mathcal{P}_{proc}$ for the probability measures in $\mathcal{M}_{type}(\mathcal{S})$ and $\mathcal{M}_{proc}(\mathcal{S})$, respectively.

The MDP induced by a pBPP. To a pBPP $\mathcal{S} = (\Gamma, \hookrightarrow, Prob)$ we associate the MDP $\mathcal{D}(\mathcal{S}) = (\mathbb{N}^\Gamma, \Gamma \cup \{\bot\}, En, \delta)$ with a fresh action $\bot \notin \Gamma$, and $En(\alpha) = \{X \in \Gamma \mid \alpha(X) \geq 1\}$ for $\varepsilon \neq \alpha \in \mathbb{N}^\Gamma$ and $En(\varepsilon) = \{\bot\}$, and $\delta(\alpha, X)(\alpha - X + \beta) = p$ whenever $\alpha(X) \geq 1$ and $X \overset{p}{\hookrightarrow} \beta$, and $\delta(\varepsilon, \bot)(\varepsilon) = 1$. As in the Markov chain, the new configuration γ is obtained from α by replacing an X-process with a configuration randomly sampled according to the rules $X \overset{p}{\hookrightarrow} \beta$. But in contrast to the Markov chain the selection of X is up to a scheduler.

3 The Coverability Problem for the Markov Chain

In this section we study the *coverability problem* for the Markov chains induced by a pBPP. We say a run $\alpha_0 \alpha_1 \cdots$ of a pBPP $\mathcal{S} = (\Gamma, \hookrightarrow, Prob)$ *covers* a configuration $\phi \in \mathbb{N}^\Gamma$, if $\alpha_i \geq \phi$ holds for some $i \in \mathbb{N}$. The coverability problem asks whether it is almost surely the case that some configuration from a finite set $\{\phi_1, \ldots, \phi_n\}$ will be covered. More formally, the coverability problem is the following. Given a pBPP $\mathcal{S} = (\Gamma, \hookrightarrow, Prob)$, an initial configuration $\alpha_0 \in \mathbb{N}^\Gamma$, and finitely many configurations $\phi_1, \ldots, \phi_n$, does $\mathcal{P}_{type}(\alpha_0 \models \Diamond F) = 1$ hold,

where $F = \phi_1\uparrow \cup \ldots \cup \phi_n\uparrow$? Similarly, does $\mathcal{P}_{proc}(\alpha_0 \models \Diamond F) = 1$ hold? We will show that those two questions always have the same answer.

In Section 3.1 we show that the coverability problem is decidable. In Section 3.2 we show that the complexity of the coverability problem is nonelementary.

3.1 Decidability

For deciding the coverability problem we use the approach of [1]. The following proposition is crucial for us:

Proposition 1. *Let $\mathcal{M} = (Q, \delta)$ be a Markov chain and $F \subseteq Q$ such that $\mathcal{M}$ is globally coarse with respect to F. Let $\bar{F} = Q \setminus F$ be the complement of F and let $\widetilde{F} := \{s \in Q \mid \mathcal{P}(s \models \Diamond F) = 0\} \subseteq \bar{F}$ denote the set of states from which F is not reachable in $\mathcal{M}$. Let $s_0 \in Q$. Then we have $\mathcal{P}(s_0 \models \Diamond F) = 1$ if and only if $\mathcal{P}(s_0 \models \bar{F} \mathsf{U} \widetilde{F}) = 0$.*

Proof. Immediate from [1, Lemmas 3.7, 5.1 and 5.2]. □

In other words, Proposition 1 states that F is almost surely reached if and only if there is no path to $\widetilde{F}$ that avoids F. Proposition 1 will allow us to decide the coverability problem by computing only reachability relations in $\mathcal{M}$, ignoring the probabilities.

Recall that for a pBPP $\mathcal{S} = (\Gamma, \hookrightarrow, Prob)$, the Markov chains $\mathcal{M}_{type}(\mathcal{S})$ and $\mathcal{M}_{proc}(\mathcal{S})$ have the same structure; only the transition probabilities differ. In particular, if $F \subseteq \mathbb{N}^\Gamma$ is upward-closed, the set $\widetilde{F}$, as defined in Proposition 1, is the same for $\mathcal{M}_{type}(\mathcal{S})$ and $\mathcal{M}_{proc}(\mathcal{S})$. Moreover, we have the following proposition (full proof in Appendix A).

Proposition 2. *Let $\mathcal{S} = (\Gamma, \hookrightarrow, Prob)$ be a pBPP. Let $F \subseteq \mathbb{N}^\Gamma$ be upward-closed. Then the Markov chains $\mathcal{M}_{type}(\mathcal{S})$ and $\mathcal{M}_{proc}(\mathcal{S})$ are globally coarse with respect to F.*

Proof (sketch). The statement about $\mathcal{M}_{type}(\mathcal{S})$ follows from [1, Theorem 4.3]. For the statement about $\mathcal{M}_{proc}(\mathcal{S})$ it is crucial to argue that starting with any configuration $\alpha \in \mathbb{N}^\Gamma$ it is the case with probability 1 that every type X with $\alpha(X) \geq 1$ is eventually scheduled. Since F is upward-closed it follows that for all $\alpha, \beta \in \mathbb{N}^\Gamma$ with $\alpha \leq \beta$ we have $\mathcal{P}_{proc}(\alpha \models \Diamond F) \leq \mathcal{P}_{proc}(\beta \models \Diamond F)$. Then the statement follows from Dickson's lemma.

For an illustration of the challenge, consider the pBPP with $X \overset{1}{\hookrightarrow} XX$ and $Y \overset{1}{\hookrightarrow} YY$, and let $F = XX\uparrow$. Clearly we have $\mathcal{P}_{proc}(X \models \Diamond F) = 1$, as the X-process is scheduled immediately. Now let $\alpha_0 = XY$. Since $\alpha_0 \geq X$, the inequality claimed above implies $\mathcal{P}_{proc}(\alpha_0 \models F) = 1$. Indeed, the probability that the X-process in α_0 is *never* scheduled is at most $\frac{1}{2} \cdot \frac{2}{3} \cdot \frac{3}{4} \cdot \ldots$, which is 0. Hence $\mathcal{P}_{proc}(\alpha_0 \models \Diamond F) = 1$. □

The following proposition follows by combining Propositions 1 and 2.

Proposition 3. *Let $\mathcal{S} = (\Gamma, \hookrightarrow, Prob)$ be a pBPP. Let $F \subseteq \mathbb{N}^\Gamma$ be upward-closed. Let $\alpha_0 \in \mathbb{N}^\Gamma$. We have:*

$$\mathcal{P}_{type}(\alpha_0 \models \Diamond F) = 1 \iff \mathcal{P}_{type}(\alpha_0 \models \bar{F} \mathsf{U} \widetilde{F}) = 0$$
$$\iff \mathcal{P}_{proc}(\alpha_0 \models \bar{F} \mathsf{U} \widetilde{F}) = 0 \iff \mathcal{P}_{proc}(\alpha_0 \models \Diamond F) = 1$$

By Proposition 3 we may in the following omit the subscript from $\mathcal{P}_{type}, \mathcal{P}_{proc}, \mathcal{M}_{type}, \mathcal{M}_{proc}$ if it does not matter. We have the following theorem.

Theorem 4. *The coverability problem is decidable: given a pBPP $\mathcal{S} = (\Gamma, \hookrightarrow, Prob)$, an upward-closed set $F \subseteq \mathbb{N}^\Gamma$, and a configuration $\alpha_0 \in \mathbb{N}^\Gamma$, it is decidable whether $\mathcal{P}(\alpha_0 \models \Diamond F) = 1$ holds.*

Proof. The complement of $\widetilde{F}$ (i.e., the set of configurations from which F is reachable) is upward-closed, and its minimal elements can be computed by a straightforward fixed-point computation (this is even true for the more general model of pVASS, e.g., see [1, Remark 4.2]). By Proposition 3 it suffices to decide whether $\mathcal{P}(\alpha_0 \models \bar{F} \mathsf{U} \widetilde{F}) > 0$ holds. Define $R := \{\alpha \in \bar{F} \mid \alpha \text{ is reachable from } \alpha_0 \text{ via } \bar{F}\text{-configurations}\}$. Observe that $\mathcal{P}(\alpha_0 \models \bar{F} \mathsf{U} \widetilde{F}) > 0$ if and only if $R \cap \widetilde{F} \neq \emptyset$. We can now give a Karp-Miller-style algorithm for checking that $R \cap \widetilde{F} \neq \emptyset$: (i) Starting from α_0, build a tree of configurations reachable from α_0 via $\bar{F}$-configurations (i.e., at no stage F-configurations are added to this tree) — for example, in a breadth-first search manner — but stop expanding a leaf node α_k as soon as we discover that the branch $\alpha_0 \to \cdots \to \alpha_k$ satisfies the following: $\alpha_j \leq \alpha_k$ for some $j < k$. (ii) As soon as a node $\alpha \in \widetilde{F}$ is generated, terminate and output "yes". (iii) When the tree construction is completed without finding nodes in $\widetilde{F}$, terminate and output "no".

To prove correctness of the above algorithm, we first prove termination. To this end, it suffices to show that the constructed tree is finite. To see this, observe first that every branch in the constructed tree is of finite length. This is an immediate consequence of Dickon's lemma and our policy of terminating a leaf node α that satisfies $\alpha' \leq \alpha$, for some ancestor α' of α in this tree. Now since all branches of the tree are finite, König's lemma shows that the tree itself must be finite (since each node has finite degree).

To prove partial correctness, it suffices to show that the policy of terminating a leaf node α that satisfies $\alpha' \leq \alpha$, for some ancestor α' of α in this tree, is valid. That is, we want to show that if $R \cap \widetilde{F} \neq \emptyset$ then a witnessing vector $\gamma \in R \cap \widetilde{F}$ will be found by the algorithm. We have the following lemma whose proof is in Appendix A.

Lemma 5. *Let $\alpha_0 \in \bar{F}$ and let $\gamma \in \mathbb{N}^\Gamma$. Let $\alpha_0 \to \alpha_1 \to \ldots \to \alpha_k$ be a shortest path in $\mathcal{M}(\mathcal{S})$ such that $\alpha_0, \ldots, \alpha_k \in \bar{F}$ and $\alpha_k \leq \gamma$. Then for all i, j with $0 \leq i < j \leq k$ we have $\alpha_i \not\leq \alpha_j$.*

Let $R \cap \widetilde{F} \neq \emptyset$ and let $\gamma \in \mathbb{N}^\Gamma$ be a minimal element of $R \cap \widetilde{F}$. By Lemma 5 our algorithm does not prune any shortest path from α_0 to γ. Hence it outputs "yes". □

3.2 Nonelementary Lower Bound

We have the following lower-bound result:

Theorem 6. *The complexity of the coverability problem is nonelementary.*

The proof is technically involved.

Proof (sketch). We claim that there exists a nonelementary function f such that given a 2-counter machine M running in space $f(k)$, we can compute a pBPP $\mathcal{S} = (\Gamma, \hookrightarrow, Prob)$ of size $\leq k$, an upward-closed set $F \subseteq \mathbb{N}^\Gamma$ (with at most k minimal elements, described by numbers at most k), and a type $X_0 \in \Gamma$, such that $\mathcal{P}(X_0 \models \Diamond F) = 1$ holds if and only if M does not terminate. Recall that by Proposition 3 we have that $\mathcal{P}(X_0 \models \Diamond F) = 1$ is equivalent to $\mathcal{P}(X_0 \models \bar{F} \mathsf{U} \widetilde{F}) = 0$.

Since the exact values of the probabilities do not matter, it suffices to construct a (nonprobabilistic) BPP $\mathcal{S}$. Further, by adding processes that can spawn everything (and hence cannot take part in $\widetilde{F}$-configurations) one can change the condition of reaching $\widetilde{F}$ to reaching a downward closed set $G \subseteq \bar{F}$. So the problem we are reducing to is: does there exist a path in $\mathcal{S}$ that is contained in $\bar{F}$ and goes from X_0 to a downward closed set G.

By defining F suitably we can add various restrictions on the behaviour of our BPP. For example, the following example allows X to be turned into Y if and only if there is no Z present:

$$X \hookrightarrow YW \qquad W \hookrightarrow \varepsilon \qquad F = WZ\uparrow$$

Doubling the number of a given process is straightforward, and it is also possible to divide the number of a given process by two. Looking only at runs inside $\bar{F}$, the following BPP can turn all its X-processes into half as many X'-processes. (Note that more X'-processes could be spawned, but because of the monotonicity of the system, the "best" runs are those that spawn a minimal number of processes.)

$$X \hookrightarrow TP \qquad T \hookrightarrow \overline{P} \qquad P \hookrightarrow \varepsilon \qquad \overline{P} \hookrightarrow \varepsilon$$

$$P_1 \hookrightarrow \overline{P_2} \qquad \overline{P_2} \hookrightarrow P_2 \qquad P_2 \hookrightarrow \overline{P_1} \qquad \overline{P_1} \hookrightarrow P_1 X'$$

$$F = P\overline{P_1}\uparrow \cup P\overline{P_2}\uparrow \cup \overline{P}P_1\uparrow \cup \overline{P}P_2\uparrow \cup T^2\uparrow$$

$$\alpha_{init} = X^n\overline{P_1}$$

Let us explain this construction. In order to make an X-process disappear, we need to create temporary processes P and $\overline{P}$. However, these processes are incompatible, respectively, with $\overline{P_i}$ and P_i. Thus, destroying an X-process requires the process P_1 to move into $\overline{P_1}$ and then into P_2. By repeatedly destroying X-processes, this forces the creation of half as many X'-processes.

It is essential for our construction to have a loop-gadget that performs a cycle of processes $A \hookrightarrow B \hookrightarrow C \hookrightarrow A$ exactly k times ("k-loop"). By activating/disabling transitions based on the absence/presence of an A-, B- or

C-process, we can force an operation to be performed k times. For example, assuming the construction of a k-loop gadget, the following BPP doubles the number of X-processes k times:

$$X \hookrightarrow Y \qquad Y \hookrightarrow ZZ \qquad Z \hookrightarrow X$$

(rules for k-loop on $A/B/C$)

$$F = XB\uparrow \cup YC\uparrow \cup ZA\uparrow$$

For the loop to perform $A \hookrightarrow B$, all X-processes have to be turned into Y. Similarly, performing $B \hookrightarrow C \hookrightarrow A$ requires the Y-processes to be turned into Z, then into X. Thus, in order to perform one iteration of the loop, one needs to double the number of X-processes.

To implement such a loop we need two more gadgets: one for creating k processes, and one for consuming k processes. By turning a created process into a consumed process on at a time, we obtain the required cycle. Here is an example:

$$I \hookrightarrow A \qquad A \hookrightarrow B \qquad B \hookrightarrow C \qquad C \hookrightarrow \varepsilon$$

$$\overline{A} \hookrightarrow \overline{B} \qquad \overline{B} \hookrightarrow \overline{C}F \qquad \overline{C} \hookrightarrow A$$

(rules for a gadget to consume k processes F)

(rules for a gadget to spawn k processes I)

$$F = A\overline{A}\uparrow \cup B\overline{B}\uparrow \cup C\overline{C}\uparrow \cup AA\uparrow \cup BB\uparrow \cup CC\uparrow$$

$$\alpha_{init} = \overline{A}$$

By combining a k-loop with a multiplier or a divider, we can spawn or consume 2^k processes. This allows us to create a 2^k-loop. By iterating this construction, we get a BPP of exponential size (each loop requires two lower-level loops) that is able to spawn or consume $2^{2^{\cdots^k}}$ processes.

It remains to simulate our 2-counter machine M. The main idea is to spawn an initial budget b of processes, and to make sure that this number stays the same along the run. Zero-tests are easy to implement; the difficulty lies in the increments and decrements. The solution is to maintain, for each simulated counter, two pools of processes X and $\overline{X}$, such that if the counter is supposed to have value k, then we have processes $X^k\overline{X}^{b-k}$. Now, incrementing consist in turning all these processes into backup processes, except one $\overline{X}$-process. Then, we turn this process into an X-process, and return all backup process to their initial type.

Appendix B provides complete details of the proofs, including graphical representations of the processes involved. □

4 The Coverability Problem for the MDP

In the following we investigate the controlled version of the pBPP model. Recall from Section 2 that a pBPP $\mathcal{S} = (\Gamma, \hookrightarrow, Prob)$ induces an MDP $\mathcal{D}(\mathcal{S})$ where in a configuration $\varepsilon \neq \alpha \in \mathbb{N}^\Gamma$ a scheduler σ picks a type X with $\alpha(X) \geq 1$. The successor configuration is then obtained randomly from α according to the rules in $\mathcal{S}$ with X on the left-hand side.

We investigate (variants of) the decision problem that asks, given $\alpha_0 \in \mathbb{N}^\Gamma$ and an upward-closed set $F \subseteq \mathbb{N}^\Gamma$, whether $\mathcal{P}_\sigma(\alpha_0 \models \Diamond F) = 1$ holds for some scheduler (or for all schedulers, respectively).

4.1 The Existential Problem

In this section we consider the scenario where we ask for a scheduler that makes the system reach an upward-closed set with probability 1. We prove the following theorem:

Theorem 7. *Given a pBPP $\mathcal{S} = (\Gamma, \hookrightarrow, Prob)$ and a configuration $\alpha_0 \in \mathbb{N}^\Gamma$ and an upward-closed set $F \subseteq \mathbb{N}^\Gamma$, it is decidable whether there exists a scheduler σ with $\mathcal{P}_\sigma(\alpha_0 \models \Diamond F) = 1$. If such a scheduler exists, one can compute a deterministic and memoryless scheduler σ with $\mathcal{P}_\sigma(\alpha_0 \models \Diamond F) = 1$.*

Proof (sketch). The proof (in Appendix C) is relatively long. The idea is to abstract the MDP $\mathcal{D}(\mathcal{S})$ (with $\mathbb{N}^\Gamma$ as state space) to an "equivalent" finite-state MDP. The state space of the finite-state MDP is $Q := \{0, 1, \ldots, K\}^\Gamma \subseteq \mathbb{N}^\Gamma$, where K is the largest number that appears in the minimal elements of F. For finite-state MDPs, reachability with probability 1 can be decided in polynomial time, and an optimal deterministic and memoryless scheduler can be synthesized.

When setting up the finite-state MDP, special care needs to be taken of transitions that would lead from Q to a configuration α outside of Q, i.e., $\alpha \in \mathbb{N}^\Gamma \setminus Q$. Those transitions are redirected to a probability distribution on T_α with $T_\alpha \subseteq Q$, so that each configuration in T_α is "equivalent" to some configuration $\beta \in \mathbb{N}^\Gamma$ that could be reached from α in the infinite-state MDP $\mathcal{D}(\mathcal{S})$, if the scheduler follows a particular optimal strategy in $\mathcal{D}(\mathcal{S})$. (One needs to show that indeed with probability 1 such a β is reached in the infinite-state MDP, if the scheduler acts according to this strategy.) This optimal strategy is based on the observation that whenever in configuration $\beta \in \mathbb{N}^\Gamma$ with $\beta(X) > K$ for some X, then type X can be scheduled. This is without risk, because after scheduling X, at least K processes of type X remain, which is enough by the definition of K. The benefit of scheduling such X is that processes appearing on the right-hand side of X-rules may be generated, possibly helping to reach F. For computing T_α, we rely on decision procedures for the reachability problem in Petri nets, which prohibits us from giving an upper complexity bound. □

4.2 The Universal Problem

In this section we consider the scheduler as adversarial in the sense that it tries to avoid the upward-closed set F. We say "the scheduler wins" if it avoids F forever.

We ask if the scheduler can win with positive probability: given α_0 and F, do we have $\mathcal{P}_\sigma(\alpha_0 \models \Diamond F) = 1$ for all schedulers σ? For the question to make sense, we need to rephrase it, as we show now. Consider the pBPP $\mathcal{S} = (\Gamma, \hookrightarrow, \mathit{Prob})$ with $\Gamma = \{X, Y\}$ and the rules $X \overset{1}{\hookrightarrow} XX$ and $Y \overset{1}{\hookrightarrow} YY$. Let $F = XX\uparrow$. If $\alpha_0 = X$, then, clearly, we have $\mathcal{P}_\sigma(\alpha_0 \models \Diamond F) = 1$ for all schedulers σ. However, if $\alpha_0 = XY$, then there is a scheduler σ with $\mathcal{P}_\sigma(\alpha_0 \models \Diamond F) = 0$: take the scheduler σ that always schedules Y and never X. Such a scheduler is intuitively *unfair*. If an operating system acts as a scheduler, a minimum requirement would be that waiting processes are scheduled eventually.

We call a run $\alpha_0 X_1 \alpha_1 X_2 \ldots$ in the MDP $\mathcal{D}(\mathcal{S})$ *fair* if for all $i \geq 0$ and all $X \in \Gamma$ with $\alpha_i(X) \geq 1$ we have $X = X_j$ for some $j > i$. We call a scheduler σ *classically fair* if it produces only fair runs.

Example 8. Consider the pBPP with $X \overset{1}{\hookrightarrow} Y$ and $Y \overset{0.5}{\hookrightarrow} Y$ and $Y \overset{0.5}{\hookrightarrow} X$. Let $F = YY\uparrow$. Let $\alpha_0 = XX$. In configuration $\alpha = XY$ the scheduler has to choose between two options: It can pick X, resulting in the successor configuration $YY \in F$, which is a "loss" for the scheduler. Alternatively, it picks Y, which results in α or α_0, each with probability 0.5. If it results in α, nothing has changed; if it results in α_0, we say a "a round is completed". Consider the scheduler σ that acts as follows. When in configuration $\alpha = XY$ and in the ith round, it picks Y until either the next round (the $(i+1)$st round) is completed or Y has been picked i times in this round. In the latter case it picks X and thus loses. Clearly, σ is classically fair (provided that it behaves in a classically fair way after it loses, for instances using round-robin). The probability of losing in the ith round is 2^{-i}. Hence the probability of losing is $\mathcal{P}_\sigma(\alpha_0 \models \Diamond F) = 1 - \prod_{i=1}^{\infty}(1 - 2^{-i}) < 1$. (For this inequality, recall that for a sequence $(a_i)_{i\in\mathbb{N}}$ with $a_i \in (0,1)$ we have $\prod_{i\in\mathbb{N}}(1-a_i) = 0$ if and only if the series $\sum_{i\in\mathbb{N}} a_i$ diverges.) One can argue along these lines that any classically fair scheduler needs to play longer and longer rounds in order to win with positive probability. In particular, such schedulers need infinite memory.

It is hardly conceivable that an operating system would "consider" such schedulers. Note that the pBPP from the previous example has a finite state space.

In the probabilistic context, a commonly used alternative notion is *probabilistic fairness*, see e.g. [10, 17] or [4] for an overview (the term *probabilistic fairness* is used differently in [4]). We call a scheduler σ *probabilistically fair* if with probability 1 it produces a fair run.

Example 9. For the pBPP from the previous example, consider the scheduler σ that picks Y until the round is completed. Then $\mathcal{P}_\sigma(\alpha \models \Diamond F) = 0$ and σ is probabilistically fair.

The following example shows that probabilistic fairness for pBPPs can be unstable with respect to perturbations in the probabilities.

Example 10. Consider a pBPP with

$$X \overset{1}{\hookrightarrow} Y \qquad Y \overset{1}{\hookrightarrow} XZ \qquad Z \overset{p}{\hookrightarrow} ZZ \qquad Z \overset{1-p}{\hookrightarrow} \varepsilon \qquad \text{for some } p \in (0,1)$$

and $F = YZ\uparrow$ and $\alpha_0 = XZ$.

Let $p \leq 0.5$. Then, by an argument on the "gambler's ruin problem" (see e.g. [8, Chapter XIV]), with probability 1 each Z-process produces only finitely many other Z-processes in its "subderivation tree". Consider the scheduler σ that picks Z as long as there is a Z-process. With probability 1 it creates a run of the following form:

$$(XZ)\cdots(X)(Y)(XZ)\cdots(X)(Y)(XZ)\cdots(X)(Y)(XZ)\ldots$$

Such runs are fair, so σ is probabilistically fair and wins with probability 1.

Let $p > 0.5$. Then, by the same random-walk argument, with probability 1 some Z-process (i.e., at least one of the Z-processes created by Y) produces infinitely many other Z-processes. So any probabilistically fair scheduler σ produces, with probability 1, a Y-process before all Z-processes are gone, and thus loses.

We conclude that a probabilistically fair scheduler σ with $\mathcal{P}_\sigma(\alpha_0 \models \Diamond F) < 1$ exists if and only if $p \leq 0.5$.

The example suggests that deciding whether there exists a probabilistically fair scheduler σ with $\mathcal{P}_\sigma(\alpha_0 \models \Diamond F) < 1$ requires arguments on (in general) multidimensional random walks. In addition, the example shows that probabilistic fairness is not a robust notion when the exact probabilities are not known.

We aim at solving those problems by considering a stronger notion of fair runs: Let $k \in \mathbb{N}$. We call a run $\alpha_0 X_1 \alpha_1 X_2 \ldots$ *k-fair* if for all $i \geq 0$ and all $X \in \Gamma$ with $\alpha_i(X) \geq 1$ we have that $X \in \{X_{i+1}, X_{i+2}, \ldots, X_k\}$. In words, if $\alpha_i(X) \geq 1$, the type X has to be scheduled within time k. We call a scheduler *k-fair* if it produces only k-fair runs.

Theorem 11. *Given a pBPP $\mathcal{S} = (\Gamma, \hookrightarrow, Prob)$, an upward-closed set F, a number $k \in \mathbb{N}$, and a configuration $\alpha_0 \in \mathbb{N}^\Gamma$, it is decidable whether for all k-fair schedulers σ we have $\mathcal{P}_\sigma(\alpha_0 \models \Diamond F) = 1$.*

The proof is inspired by proofs in [16], and combines new insights with the technique of Theorem 4, see Appendix C. We remark that the proof shows that the exact values of the positive probabilities do not matter.

5 Q-States Target Sets

In this section, we provide a sensible restriction of input target sets which yields polynomial-time solvability of our problems. Let $Q = \{X_1, \ldots, X_n\} \subseteq \Gamma$. The *$Q$-states set* is the upward-closed set $F = X_1\uparrow \cup \ldots \cup X_n\uparrow$. There are two reasons to consider Q-states target sets. Firstly, Q-states target sets are sufficiently expressive to capture common examples in the literature of distributed protocols, e.g., freedom from deadlock and resource starvation (standard examples include the dining philosopher problem in which case *at least one* philosopher must eat). Secondly, Q-states target sets have been considered in the literature

of Petri nets: e.g., in [1][1] the authors showed that qualitative reachability for probabilistic Vector Addition Systems with States with Q-states target sets becomes decidable whereas the same problem is undecidable with upward-closed target sets.

Theorem 12. *Let $\mathcal{S} = (\Gamma, \hookrightarrow, Prob)$ be a pBPP. Let $Q \subseteq \Gamma$ represent an upward-closed set $F \subseteq \mathbb{N}^\Gamma$. Let $\alpha_0 \in \mathbb{N}^\Gamma$ and $k \geq |\Gamma|$.*

(a) The coverability problem with Q-states target sets is solvable in polynomial time; i.e., we can decide in polynomial time whether $\mathcal{P}(\alpha_0 \models \Diamond F) = 1$ holds.

(b) We have:

$$\begin{aligned} & \mathcal{P}(\alpha_0 \models \Diamond F) = 1 \\ \iff\ & \mathcal{P}_\sigma(\alpha_0 \models \Diamond F) = 1 \text{ holds for some scheduler } \sigma \\ \iff\ & \mathcal{P}_\sigma(\alpha_0 \models \Diamond F) = 1 \text{ holds for all } k\text{-fair schedulers } \sigma. \end{aligned}$$

As a consequence of part (a), the existential and the k-fair universal problem are decidable in polynomial time.

Proof. Denote by $Q' \subseteq \Gamma$ the set of types $X \in \Gamma$ such that there are $\ell \in \mathbb{N}$, a path $\alpha_0, \ldots, \alpha_\ell$ in the Markov chain $\mathcal{M}(\mathcal{S})$, and a type $Y \in Q$ such that $\alpha_0 = X$ and $\beta = \alpha_\ell \geq Y$. Clearly we have $Q \subseteq Q' \subseteq \Gamma$, and Q' can be computed in polynomial time.

In the following, view $\mathcal{S}$ as a context-free grammar with empty terminal set (ignore the probabilities, and put the symbols on the right-hand sides in an arbitrary order). Remove from $\mathcal{S}$ all rules of the form: (i) $X \hookrightarrow \alpha$ where $X \in Q$ or $\alpha(Y) \geq 1$ for some $Y \in Q$, and (ii) $X \hookrightarrow \alpha$ where $X \in \Gamma \setminus Q'$. Furthermore, add rules $X \hookrightarrow \varepsilon$ where $X \in \Gamma \setminus Q'$. Check (in polynomial time) whether in the grammar the empty word ε is produced by α_0.

We have that ε is produced by α_0 if and only if $\mathcal{P}(\alpha_0 \models \Diamond F) < 1$. This follows from Proposition 3, as the complement of $\widetilde{F}$ is the Q'-states set. Hence part (a) of the theorem follows.

For part (b), let $\mathcal{P}(\alpha_0 \models \Diamond F) < 1$. By part (a) we have that ε is produced by α_0. Then for all schedulers σ we have $\mathcal{P}_\sigma(\alpha_0 \models \Diamond F) < 1$. Trivially, as a special case, this holds for some k-fair scheduler. (Note that k-fair schedulers exist, as $k \geq |\Gamma|$.)

Conversely, let $\mathcal{P}(\alpha_0 \models \Diamond F) = 1$. By part (a) we have that ε is not produced by α_0. Then, no matter what the scheduler does, the set F remains reachable. So all k-fair schedulers will, with probability 1, hit F eventually. □

6 Semilinear Target Sets

In this section, we prove that the qualitative reachability problems that we considered in the previous sections become undecidable when we extend upward-closed to semilinear target sets.

[1] Our definition seems different from [1], but equivalent from standard embedding of Vector Addition Systems with States to Petri Nets.

Theorem 13. *Let $\mathcal{S} = (\Gamma, \hookrightarrow, Prob)$ be a pBPP. Let $F \subseteq \mathbb{N}^\Gamma$ be a semilinear set. Let $\alpha_0 \in \mathbb{N}^\Gamma$. The following problems are undecidable:*

(a) Does $\mathcal{P}(\alpha_0 \models \Diamond F) = 1$ hold?
(b) Does $\mathcal{P}_\sigma(\alpha_0 \models \Diamond F) = 1$ hold for all 7-fair schedulers σ?
(c) Does $\mathcal{P}_\sigma(\alpha_0 \models \Diamond F) = 1$ hold for some scheduler σ?

The proofs are reductions from the control-state-reachability problem for 2-counter machines, see Appendix D.

7 Conclusions and Future Work

In this paper we have studied fundamental qualitative coverability and other reachability properties for pBPPs. For the Markov-chain model, the coverability problem for pBPPs is decidable, which is in contrast to general pVASSs. We have also shown a nonelementary lower complexity bound. For the MDP model, we have proved decidability of the existential and the k-fair version of the universal coverability problem. The decision algorithms for the MDP model are not (known to be) elementary, as they rely on Petri-net reachability and a Karp-Miller-style construction, respectively. It is an open question whether there exist elementary algorithms. Another open question is whether the universal MDP problem without any fairness constraints is decidable.

We have given examples of problems where the answer depends on the exact probabilities in the pBPP. This is also true for the reachability problem for finite sets: Given a pBPP and $\alpha_0 \in \mathbb{N}^\Gamma$ and a *finite* set $F \subseteq \mathbb{N}^\Gamma$, the reachability problem for finite sets asks whether we have $\mathcal{P}(\alpha_0 \models \Diamond F) = 1$ in the Markov chain $\mathcal{M}(\mathcal{S})$. Similarly as in Example 10 the answer may depend on the exact probabilities: consider the pBPP with $X \overset{p}{\hookrightarrow} XX$ and $X \overset{1-p}{\hookrightarrow} \varepsilon$, and let $\alpha_0 = XX$ and $F = \{X\}$. Then we have $\mathcal{P}(\alpha_0 \models \Diamond F) = 1$ if and only if $p \leq 1/2$. The same is true in both the existential and the universal MDP version of this problem. Decidability of all these problems is open, but clearly decision algorithms would have to use techniques that are very different from ours, such as analyses of multidimensional random walks.

On a more conceptual level we remark that the problems studied in this paper are qualitative in two senses: (a) we ask whether certain events happen with probability 1 (rather than > 0.5 etc.); and (b) the exact probabilities in the rules of the given pBPP do not matter. Even if the system is nondeterministic and not probabilistic, properties (a) and (b) allow for an interpretation of our results in terms of nondeterministic BPPs, where the nondeterminism is constrained by the laws of probability, thus imposing a special but natural kind of fairness. It would be interesting to explore this kind of "weak" notion of probability for other (infinite-state) systems.

Acknowledgment Anthony W. Lin did this work when he was an EPSRC research fellow at Oxford University supported by grant number EP/H026878/1.

References

1. P.A. Abdulla, N. Ben Henda, and R. Mayr. Decisive Markov chains. *Logical Methods in Computer Science*, 3(4:7), 2007.
2. T. Brázdil, J. Esparza, S. Kiefer, and M. Luttenberger. Space-efficient scheduling of stochastically generated tasks. *Information and Computation*, 210:87–110, 2012.
3. K. Chatterjee and M. Henzinger. Faster and dynamic algorithms for maximal end-component decomposition and related graph problems in probabilistic verification. In *Proceedings of SODA*, pages 1318–1336, 2011.
4. L. de Alfaro. From fairness to chance. *Electronic Notes in Theoretical Computer Science*, 22:55–87, 1999.
5. N. Dershowitz, D.N. Jayasimha, and S. Park. Bounded fairness. In *Verification: Theory and Practice*, pages 304–317, 2003.
6. J. Esparza. Petri nets, commutative context-free grammars, and basic parallel processes. *Fundam. Inform.*, 31(1):13–25, 1997.
7. J. Esparza and A. Kiehn. On the model checking problem for branching time logics and basic parallel processes. In *Proceedings of CAV*, number 939 in LNCS, pages 353–366, 1995.
8. W. Feller. *An introduction to probability theory and its applications*, volume I. John Wiley & Sons, 1968.
9. S.B. Fröschle, P. Jančar, S. Lasota, and Z. Sawa. Non-interleaving bisimulation equivalences on basic parallel processes. *Inf. Comput.*, 208(1):42–62, 2010.
10. S. Hart, M. Sharir, and A. Pnueli. Termination of probabilistic concurrent programs. *ACM Transactions on Programming Languages and Systems*, 5(3):356–380, 1983.
11. Y. Hirshfeld, M. Jerrum, and F. Moller. A polynomial algorithm for deciding bisimilarity of normed context-free processes. *Theor. Comput. Sci.*, 158(1&2):143–159, 1996.
12. H. Hüttel, N. Kobayashi, and T. Suto. Undecidable equivalences for basic parallel processes. *Inf. Comput.*, 207(7):812–829, 2009.
13. P. Jančar. Strong bisimilarity on basic parallel processes is PSPACE-complete. In *Proceedings of LICS*, pages 218–227, 2003.
14. N.A. Lynch, I. Saias, and R. Segala. Proving time bounds for randomized distributed algorithms. In *PODC*, pages 314–323, 1994.
15. G. Norman. Analysing randomized distributed algorithms. In *Validation of Stochastic Systems*, pages 384–418, 2004.
16. J.-F. Raskin, M. Samuelides, and L. Van Begin. Games for counting abstractions. *Electronic Notes in Theoretical Computer Science*, 128(6):69–85, 2005.
17. M.Y. Vardi. Automatic verification of probabilistic concurrent finite state programs. In *Proceedings of FOCS*, pages 327–338, 1985.

A Proofs of Section 3

Proposition 2. *Let $\mathcal{S} = (\Gamma, \hookrightarrow, Prob)$ be a pBPP. Let $F \subseteq \mathbb{N}^\Gamma$ be upward-closed. Then the Markov chains $\mathcal{M}_{type}(\mathcal{S})$ and $\mathcal{M}_{proc}(\mathcal{S})$ are globally coarse with respect to F.*

Proof. The statement about $\mathcal{M}_{type}(\mathcal{S})$ follows from [1, Theorem 4.3]. For the statement about $\mathcal{M}_{proc}(\mathcal{S})$ define

$$Min = \{\alpha \in \mathbb{N}^\Gamma \mid \mathcal{P}_{proc}(\alpha \models \Diamond F) > 0 \text{ and for all } \alpha' < \alpha : \mathcal{P}_{proc}(\alpha' \models \Diamond F) = 0\}.$$

Note that *Min* is finite (this follows from Dickson's lemma). Define $c := \min_{\alpha \in Min} \mathcal{P}_{proc}(\alpha \models \Diamond F)$. Let $\gamma \in \mathbb{N}^\Gamma$ with $\mathcal{P}_{proc}(\gamma \models \Diamond F) > 0$. We prove the proposition by showing $\mathcal{P}_{proc}(\gamma \models \Diamond F) \geq c$.

Take $\Gamma_\bullet := \Gamma \cup \Gamma'$ where $\Gamma' = \{X' \mid X \in \Gamma\}$ is a copy of Γ. Similarly, we clone the rules so that we get $\hookrightarrow_\bullet \subseteq (\Gamma \times \mathbb{N}^\Gamma) \cup (\Gamma' \times \mathbb{N}^{\Gamma'})$ and define $Prob_\bullet$ in the obvious way. Let $\mathcal{S}_\bullet = (\Gamma_\bullet, \hookrightarrow_\bullet, Prob_\bullet)$. Let $\mathcal{P}_\bullet$ denote the probability measure of $\mathcal{M}_{proc}(\mathcal{S}_\bullet)$.

Partition γ in $\gamma = \alpha + \beta$ where $\alpha \in Min$, and let $\beta' \in \mathbb{N}^{\Gamma'}$ be a clone of β. We have:

$$\begin{aligned}
\mathcal{P}_{proc}(\gamma \models \Diamond F) &= \mathcal{P}_\bullet(\gamma \models \Diamond F) && \text{definition of } \mathcal{P}_\bullet \\
&\geq \mathcal{P}_\bullet(\alpha + \beta' \models \Diamond F) && \text{as } F \subseteq \mathbb{N}^\Gamma \\
&= \mathcal{P}_\bullet(\alpha \models \Diamond F) && \text{see below} \\
&= \mathcal{P}_{proc}(\alpha \models \Diamond F) && \text{definition of } \mathcal{P}_\bullet \\
&\geq c && \text{by definition of } c
\end{aligned}$$

To show the equality $\mathcal{P}_\bullet(\alpha + \beta' \models \Diamond F) = \mathcal{P}_\bullet(\alpha \models \Diamond F)$ we show that as long as there are Γ-processes originating from $\alpha + \beta'$ (i.e., processes originating from α), they are eventually scheduled with probability 1. In fact, let $\varepsilon \neq \alpha \in \mathbb{N}^\Gamma$ and $\beta' \in \mathbb{N}^{\Gamma'}$ be arbitrary. Let $z = \max_{X \hookrightarrow \delta} |\delta|_{proc}$ be a bound on the number of processes that can be created per step. Let $a := |\alpha|_{proc}$ and $b := |\beta'|_{proc}$. It suffices to show that the probability that only Γ'-processes are scheduled is 0. This probability is at most

$$\frac{b}{a+b} \cdot \frac{b+z}{a+b+z} \cdot \frac{b+2z}{a+b+2z} \cdot \ldots$$

Recall that for a sequence $(a_i)_{i \in \mathbb{N}}$ with $a_i \in (0,1)$ we have $\prod_{i \in \mathbb{N}}(1 - a_i) = 0$ if and only if the series $\sum_{i \in \mathbb{N}} a_i$ diverges. It follows that the infinite product above is 0. □

Lemma 5. *Let $\alpha_0 \in \bar{F}$ and let $\gamma \in \mathbb{N}^\Gamma$. Let $\alpha_0 \to \alpha_1 \to \ldots \to \alpha_k$ be a shortest path in $\mathcal{M}(\mathcal{S})$ such that $\alpha_0, \ldots, \alpha_k \in \bar{F}$ and $\alpha_k \leq \gamma$. Then for all i, j with $0 \leq i < j \leq k$ we have $\alpha_i \not\leq \alpha_j$.*

Proof. For all $\ell \in \{0, \ldots, k-1\}$ let $X_\ell \hookrightarrow \delta_\ell$ be a rule with $\alpha_\ell(X_\ell) \geq 1$ and $\alpha_\ell - X_\ell + \delta_\ell = \alpha_{\ell+1}$. Assume for a contradiction that $i < j$ with $\alpha_i \leq \alpha_j$. For all $\ell \in \{j, \ldots, k-1\}$ define $\alpha'_\ell \in \mathbb{N}^\Gamma$ and $\beta_\ell \in \mathbb{N}^\Gamma$ so that $\alpha_\ell = \alpha'_\ell + \beta_\ell$ and $\alpha'_j = \alpha_i$ and

- $\alpha'_\ell(X_\ell) \geq 1$ and $\alpha'_\ell - X_\ell + \delta_\ell = \alpha'_{\ell+1}$ or
- $\beta_\ell(X_\ell) \geq 1$ and $\beta_\ell - X_\ell + \delta_\ell = \beta_{\ell+1}$.

As $\bar{F}$ is downward-closed, we have $\alpha'_\ell \in \bar{F}$ for all $\ell \in \{j, \ldots, k\}$. It follows that

$$\alpha_0 \to \alpha_1 \to \ldots \alpha_{i-1} \to \alpha'_j \to \alpha'_{j+1} \to \ldots \to \alpha'_k \leq \gamma$$

is, after removing repetitions, a path via $\bar{F}$-states. As $i < j$, the path is shorter than the path $\alpha_0 \to \alpha_1 \to \ldots \to \alpha_k$, so we have obtained the desired contradiction. □

B Proof of the Lower Complexity Bound

Theorem 6. *The complexity of the coverability problem is nonelementary.*

We claim that there exists a nonelementary function f such that given a Turing machine $\mathcal{M}$ running in space $f(k)$, we can build a pBPP $\mathcal{S} = (\Gamma, \hookrightarrow, Prob)$ of size k, and an upward closed set $F \subseteq \mathbb{N}^\Gamma$ such that $\mathcal{P}(s_0 \models \bar{F} \cup \widetilde{F}) = 0$ if and only if $\mathcal{M}$ doesn't terminate.

First, let us mention that we can change $\widetilde{F}$ by any downward closed subset of $\bar{F}$. Assume for example that we wish to reach $\bar{G}$ for some upward closed set G. This can be done by adding new processes T and T_2, replacing F by $F_2 = T^+F \cup T_2^+G$ and adding the following transitions:

$$T \hookrightarrow T_2 \qquad T_2 \hookrightarrow \varepsilon \qquad \forall X \in \Gamma.\, T \hookrightarrow TX$$

Then, if there is a way to reach $s \in \bar{G}$ by staying into $\bar{F}$ in the original net, this means you can reach Ts by staying into $T\bar{F}$ in the modified net. Then, you can go into $T_2 s$, which is not in $T_2 G$, and then to s, that doesn't contain either T or T_2, which means that is in $\widetilde{F_2}$ given that nothing can spawn these processes anymore. Reciprocally, if you can find a way to reach $\widetilde{F_2}$ in the modified net, this means you have been able to successfully consume T (otherwise you could spawn any process), which means that there was a path in $T\bar{F}$ reaching a configuration $T_2 s$ with $s \in \bar{G}$. This path could include spurious process spawns from the process T, but by monotony, we can remove these spawns, and get a path of the original net in $\bar{F}$ reaching a configuration $s \in \bar{G}$.

We ignore the probability part, as it doesn't matter for our reachability question and we call a BPP given with a set $\bar{F}$ a *constrained BPP*, in which we consider only paths that stay inside $\bar{F}$.

Now, we build a constrained BPP that can simulate a Minsky machine. In order to do that, let us remark that, instead of defining explicitly the set F, one

can list contraints that will define the allowed set F as the union of the sets F implied by each constraint. The two basic types of constraints that we will use are:

- Processes X and Y are incompatible (i.e. if X is present, Y cannot be, and vice-versa). This is associated to $XY\uparrow$.
- Process X is unique (you can't have two copies of it). This is associated to $XX\uparrow$.

We will allow ourselves to use more complex restrictions, that can be encoded in this system by adding extra processes. These are:

- Process X prevents rule $Y \hookrightarrow u$ to be fired. This is done by adding a dummy process T and the following rules:

$$Y \hookrightarrow Tu \qquad T \hookrightarrow \varepsilon \qquad F = TX\uparrow$$

- Given a subpart of a BPP $\mathcal{N}$ (that is, a set of process types, and a set of rules refering only to these process types), with a downward closed set of initial configurations I and a downward closed set of final configurations F, the subpart $\mathcal{N}$ has an atomic behaviour: rules that don't belong to $\mathcal{N}$ can't be used if $\mathcal{N}$ is not in I or F. This is done by replacing each rule $Y \hookrightarrow u$ outside $\mathcal{N}$ by $Y \hookrightarrow uT$ and $T \hookrightarrow \varepsilon$ with the added constraint $(\overline{I \cup F})T\uparrow$.

We will use Petri Net-style depictions of BPPs, where process types are called "places" and represented as circles while processes are called tokens and are represented as bullets in their associated circle. A transition turning a process X into processes $Y_1...Y_k$ is represented as an arrow linking the place X to the places $Y_1...Y_k$. Moreover, we represent unique processes (places that can contain only one token) as squares instead of circles.

B.1 Consumers, Producers and Counters

We look at three specific kind of constrained BPP:

- A k-producer is a constrained BPP $\mathcal{N}$ with an initial configuration s_i, one data place X and a downward closet subset of final configuration F such that if $s_f \in F$ is reachable from s_i, then $s_f(X) = k$.
- A k-consumer is a constrained BPP $\mathcal{N}$ with an initial configuration s_i, one data place X and a downward closet subset of final configuration F such that for every $p \in \mathbb{N}$, $\mathcal{N}$ can reach F from $s_i + X^p$ if and only $p \leq k$.
- A k-loop is a constrained BPP $\mathcal{N}$ with three disjoint upward-closed set of configurations A, B, C, an initial configuration s_i and a downward-closed set of final configurations F such that for every run that goes from s_i to F, the net always stay in $A \cup B \cup C$ and cycles through these three sets, in the order A, B, C exactly k times.

Intuitively, a k-producer is a gadget that forces the appearance of at least k tokens. This allows to force an operation to run more than k times (by running the producer, allowing the operation to make exactly one token disappear, then require that all tokens have disappeared). Symmetrically, a k-consumer is a gadget that is able to consume up to k tokens. This allows to restrain an operation to run up to k times (by making it generate such a token, then requiring these tokens to have disappeared). Finally, a k-loop is a gadget that has a controlled cyclic behavior which occurs exactly k times. By syncing it with another gadget, this will allow to make an operation run exactly k times (it is basically a combination of a producer and a consumer).

In all the following lemmas, note that the number of constraints is polynomial in the number of places.

Lemma 14. *There exists a constant α such that given a k-loop with n places, one can build a 2^k-producer with $n + \alpha$ places.*

Proof. This producer will be of the following form:

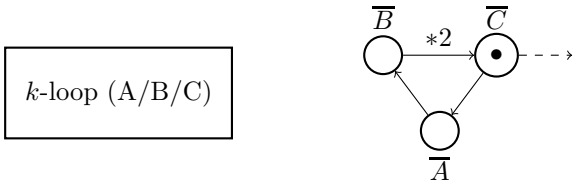

Definition of the Producer and Constraints:

- The place labelled by $\overline{C}$ is the final place that will contain the required number of tokens in the final configuration.
- The net is in its final configuration if the loop is in its final configuration and there is no more tokens in places labelled by $\overline{A}$, $\overline{B}$.
- Places labelled by $\overline{A}$ (resp. $\overline{B}$, $\overline{C}$) are incompatible with the loop being in configurations inside A (resp. B, C).

In order to put the net in its final configuration, the loop must perform k cycles on A-B-C. This means that tokens in the places $\overline{A}$, $\overline{B}$ and $\overline{C}$ must simultaneously move. Moreover, every time the loop performs a cycle, the number of tokens is doubled. This means that at the end, 2^k tokens will be in the final place $\overline{C}$.

□

Lemma 15. *There exists a constant α such that given a k-loop with n places, one can build a 2^k-consumer with $n + \alpha$ places.*

Proof. This consumer will be of the following form:

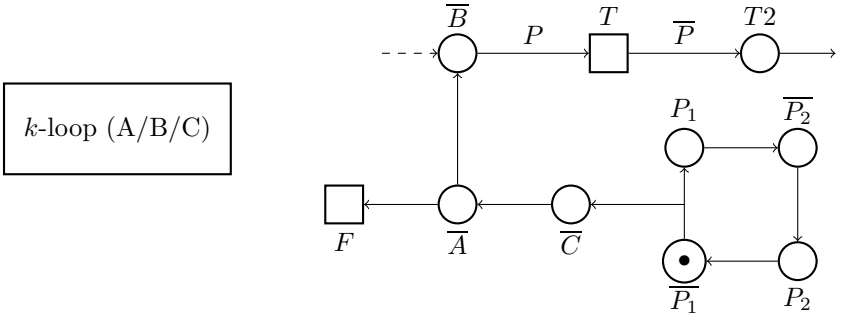

Definition of the Consumer and Constraints:

- The place labelled by $\overline{B}$ is the initial place, that will contain the initial number of tokens to consume.
- The net is in its final configuration if the loop is in its final configuration, and the only tokens in the remainder of the net are in the places labelled by $\overline{P_1}$ and F.
- Places labelled by $\overline{A}$ (resp. $\overline{B}$, $\overline{C}$) are incompatible with the loop being in configurations inside A (resp. B, C).
- The transitions labelled by P (resp. $\overline{P}$) are incompatible with the presence of tokens in the places labelled by $\overline{P_1}$ and $\overline{P_2}$ (resp. P_1 and P_2).
- The place labelled $\overline{A}$ is incompatible with the presence of a token in the places labelled by P_1, P_2, $\overline{P_2}$ and T.

Let us assume our initial configuration has n tokens in the place $\overline{B}$. We are looking to a run that empties this place. In order to do that, let us look at what happen in one step of the loop. When the loop in in configuration A, the token in the places P_1, P_2, $\overline{P_1}$ and $\overline{P_2}$ can cycle around these places. For each such cycle, two tokens can be removed from the place $\overline{B}$, and one token is created in place $\overline{C}$. Once the place $\overline{B}$ is empty, the loop can move into configuration B. There, in order to be able to move all tokens into place $\overline{A}$, we must have the place T empty, and the cycling token into place $\overline{P_1}$. Now, we can move all tokens in place $\overline{A}$, the loop in configuration C, then all tokens in place $\overline{B}$ and the loop back in configuration A. This means our net is back in its original configuration, except the loop has performed one cycle, and the number of tokens in $\overline{B}$ has been halved (rounded up). This can be done as many times as the loop allows it, with the final iteration moving one token into F instead of back into $\overline{B}$. This allows to consume up to 2^k tokens, where k is the number of iterations of the loop.

□

Lemma 16. *There exists a constant α such that given a k-consumer with n places and a k-producer with n places, one can build a k-loop with $2n + \alpha$ places.*

Proof. This loop will be of the following form:

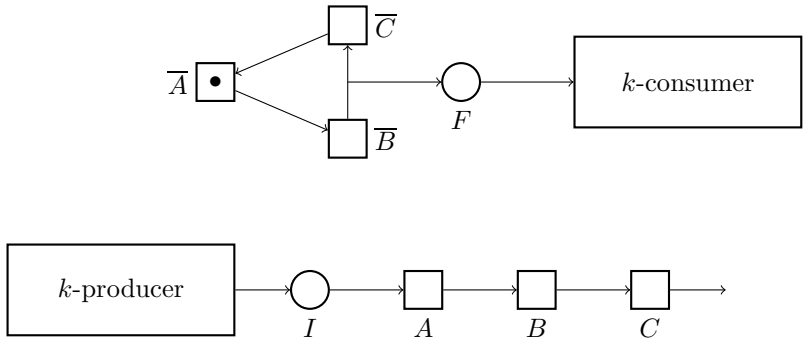

Definition of the loop and Constraints:

- The cycles of the loop are associated to the token cycling on the places $\overline{A}$, $\overline{B}$ and $\overline{C}$.
- The loop is in its final configuration if the producer and the consumer are in their final configuration, and the places I, A, B and C are empty.
- The places A, B and C are mutually exclusive.
- The place A (resp. B, C) are incompatible with the presence of a token in the place $\overline{A}$ (resp. $\overline{B}$, $\overline{C}$).

Let us look at a run going to the final configuration. In order to do that, the producer has created k tokens into place I. This means that a token has been through A, B and C at least k times, which means that the token has cycled through $\overline{A}$, $\overline{B}$ and $\overline{C}$ at least k times. Moreover, whenever such a cycle has been performed, one token has been created into F. As the consumer can consume only up to k tokens, it means there was also at most k cycles. □

B.2 Simulating a bounded counter machine

In this section, we simulate a counter machine whose counters are bounded by a constrained BPP. Our construction is made of one scheduler (see figure 1) and as many counters as the machine we want to simulate (see figure 2).

During the execution, when the scheduler is entering STEP1, the place C of each counter will contain its value, and $\overline{C}$ its complement. During steps 2 to 5, most of the tokens from C and $\overline{C}$ will be transfered respectively to B and $\overline{B}$. The places INCREMENT, DECREMENT and ZERO-TEST of each counter are called the operationnal places, and contain a token when the counter is currently performing an operation. Finally, for each transition of the machine we are simulating, we have a transition between STEP1 and STEP2 that fills for each counter the correct operationnal place.

In order to ensure our counters perform the operations requested, we have the following constraints (encoded in our upward closed set, as before):

– INIT: The producer can only run during init. He must have finished running before entering STEP1.
– STEP1: Tokens may be moved from C and $\overline{C}$ to B and $\overline{B}$. The consumer can be reset.
– STEP2: In order to enter this step, token repartition must match the token that is simultaneously appearing in the operationnal place:
  • INCREMENT: At most one token in $\overline{C}$, and none in C.
  • DECREMENT: At most one token in C, and none in $\overline{C}$.
  • ZERO-TEST: No tokens in C or B.

  Tokens in T can be consumed according to the capacity of the consumer.
– STEP3: In order to enter this step, T must be empty. Tokens may be moved freely between C and $\overline{C}$.
– STEP4: In order to enter this step, token repartition must match the token that is in the operationnal place:
  • INCREMENT: At most one token in C, and none in $\overline{C}$.
  • DECREMENT: At most one token in $\overline{C}$, and none in C.
  • ZERO-TEST: No tokens in C or B.
– STEP5: Tokens in the operational places can be deleted. Tokens in B and $\overline{B}$ can go back to C and $\overline{C}$.
– STEP6: In order to enter this step, there must be no tokens in B, $\overline{B}$ or in operational places.

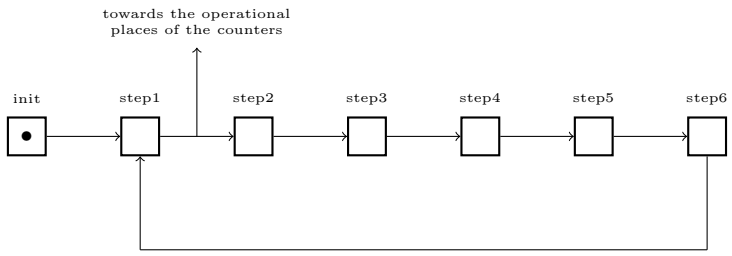

Fig. 1. The scheduler of our simulated machine

We claim that performing steps 2 to 5 moves the counter according to the instruction given by the token in INCREMENT, DECREMENT or ZERO-TEST.

The case of ZERO-TEST is simple: firing no transitions in the counter allows to move from step 2 to step 5 unimpended, and the restrictions on step 2 and 4 make sure that there must be no tokens in C before or after these operations. This means that the counter value must be zero and stays as such.

For the case of INCREMENT, the restrictions on entering step 2 means that all tokens must be moved from C to B, and all (except possibly one) must be moved from $\overline{C}$ to $\overline{B}$. However, moving a token creates a token in T, which means that

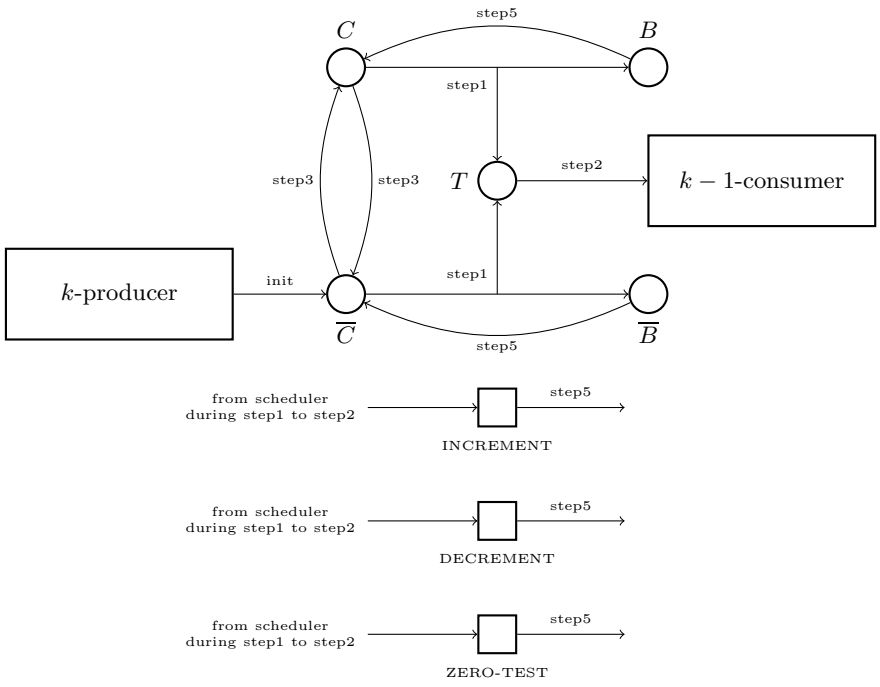

Fig. 2. A counter of our simulated machine

in order to fulfill the restrictions of step 3, at most $k-1$ tokens can be moved. Because the total number of tokens in C, B, $\overline{C}$ and $\overline{B}$ is always k, and there was no tokens in B or $\overline{B}$ before this transfer, this means that exactly $k-1$ tokens have been moved and that exactly one token remains in $\overline{C}$. Requirements of step 4 means that this token must be moved from $\overline{C}$ to C. Finally, the requirements of step 6 means that after step 5, all tokens in B and $\overline{B}$ have moved back to their original places, which means that the number of tokens in C has increased by 1.

The case of DECREMENT is symmetric.

C Proofs of Section 4

Theorem 7. *Given a pBPP $\mathcal{S} = (\Gamma, \hookrightarrow, Prob)$ and a configuration $\alpha_0 \in \mathbb{N}^\Gamma$ and an upward-closed set $F \subseteq \mathbb{N}^\Gamma$, it is decidable whether there exists a scheduler σ with $\mathcal{P}_\sigma(\alpha_0 \models \Diamond F) = 1$. If such a scheduler exists, one can compute a deterministic and memoryless scheduler σ with $\mathcal{P}_\sigma(\alpha_0 \models \Diamond F) = 1$.*

Proof. In this proof we say the scheduler *wins* if the system reaches the upward-closed set F. We also say that a configuration α is *winning* with probability 1 if there exists a scheduler σ such that $\mathcal{P}_\sigma(\alpha \models \Diamond F) = 1$. The question is whether the initial configuration α_0 is winning with probability 1.

The proof idea is to construct a finite-state MDP whose state space Q is a finite subset of $\mathbb{N}^\Gamma$. In the finite-state MDP the scheduler wins if a state from $Q \cap F$ is reached. The question whether the scheduler can win with probability 1 is decidable in polynomial time for finite-state MDPs. Moreover, for reachability with probability 1 deterministic and memoryless schedulers suffice and can be computed efficiently, see e.g. [3] or the references therein. The actions in the finite-state MDP are as in the infinite-state MDP (each type corresponds to an action), but we need to redirect transitions that would leave the finite state space.

Define the directed graph G with vertex set Γ and edges (X, Y) whenever there is a rule $X \hookrightarrow \beta$ with $\beta(Y) \geq 1$. We say that Y is a *successor* of X if (X, Y) is in the reflexive and transitive closure of the edge relation. For $X \in \Gamma$, we write $succ(X)$ for the set of successors of X. Note that $X \in succ(X)$ for all $X \in \Gamma$. Let $\phi_1, \ldots, \phi_m$ denote the minimal elements of F. Define $K := \max\{\phi_i(X) \mid 1 \leq i \leq m,\ X \in \Gamma\}$. Let $\alpha \in \mathbb{N}^\Gamma$. Define the set of *saturated* types by $Sat(\alpha) := \{X \in \Gamma \mid \forall Y \in succ(X) : \alpha(Y) \geq K\}$. Define $\lfloor \alpha \rfloor \in \mathbb{N}^\Gamma$ by $\lfloor \alpha \rfloor(X) = K$ for $X \in Sat(\alpha)$ and $\lfloor \alpha \rfloor(X) = \alpha(X)$ for $X \notin Sat(\alpha)$. Note that $Sat(\alpha) = Sat(\lfloor \alpha \rfloor)$. We make the following observation.

(1) Let $\alpha \in \mathbb{N}^\Gamma$ be winning with probability 1. Then the scheduler can win with probability 1 by *never* scheduling a type $X \in Sat(\alpha)$. Moreover, $\lfloor \alpha \rfloor$ is winning with probability 1.

For a configuration $\alpha \in \mathbb{N}^\Gamma$, define $Stable(\alpha) := \{X \in \Gamma \mid \alpha(X) \leq K\} \cup Sat(\alpha)$ and $Unstable(\alpha) := \mathbb{N}^\Gamma \setminus Stable(\alpha) = \{X \in \Gamma \mid \alpha(X) > K \text{ and } \exists Y \in$

$succ(X) : \alpha(Y) < K\}$. We call α *stable* resp. *unstable* if $Unstable(\alpha) = \emptyset$ resp. $Unstable(\alpha) \neq \emptyset$.

We define a finite-state MDP so that the scheduler can win with probability 1 if and only if it can win in the original MDP with probability 1. (In fact, we even show that the optimal winning probability stays the same.) The set of states of the finite-state MDP is

$$Q := \{\lfloor \alpha \rfloor \mid \alpha \in \mathbb{N}^\Gamma \text{ is stable}\} \subseteq \{0, \ldots, K\}^\Gamma .$$

Note that $|Q| \leq (K+1)^{|\Gamma|}$. The target states are those in F. The actions are as in the original infinite-state MDP, i.e., if $\alpha(X) \geq 1$, then scheduling X is a possible action in α. As in the infinite-state MDP, there is a special action $\perp$ enabled only in the empty configuration ε (which is losing for the scheduler except in trivial instances). If an action can lead to a state not in Q, we need to redirect those transitions to states in Q as we describe in the following. If a transition leads to a stable configuration α outside of Q, then the transition is redirected to $\lfloor \alpha \rfloor \in Q$, following Observation (1). (Also by Observation (1), the actions corresponding to types $X \in Sat(\alpha) = Sat(\lfloor \alpha \rfloor)$ could be disabled without disadvantaging the scheduler.) If a transition leads to an unstable configuration α, then this transition is redirected to a probability distribution p_α on Q so that for each $q \in Q$ we have that $p_\alpha(q)$ is the probability that in the original infinite-state MDP a configuration $\beta \in \mathbb{N}^\Gamma$ with $\lfloor \beta \rfloor = q$ is the first stable configuration reached when following a particular class of optimal strategies which we describe in the following.

The strategy class relies on the fact that in a configuration $\alpha \in \mathbb{N}^\Gamma$ with $\alpha(X) > K$ for some $X \in \Gamma$, the scheduler does not suffer a disadvantage by scheduling X: Indeed, by scheduling X, only the X-component of the configuration can decrease, and if it decreases, it decreases by at most 1; so we have $\alpha'(X) \geq K$ also for the successor configuration α'. As those X-processes in excess of K can only be useful for producing types that are successors of X, one can schedule them freely and at any time. We call a strategy *cautious* if it behaves in the following way while the current configuration $\alpha \in \mathbb{N}^\Gamma$ is unstable:

> Let $X_1, \ldots, X_k$ with $k \leq |\Gamma|$ be the shortest path (where ties are resolved in an arbitrary but deterministic way) in the graph G from the beginning of this proof such that $\alpha(X_1) > K$ and $\alpha(X_k) < K$ and $\alpha(X_i) = K$ for $2 \leq i \leq k-1$. Schedule X_1.

We claim that, with probability 1, a stable configuration will eventually be reached if a cautious strategy is followed. To see that, consider an unstable configuration α and let $X_1 \in \Gamma$ be scheduled, i.e., $X_1, \ldots, X_k$ with $k \leq |\Gamma|$ is the shortest path in the graph G such that $\alpha(X_1) > K$, and $\alpha(X_k) < K$ and $\alpha(X_i) = K$ for $2 \leq i \leq k-1$. By the definition of G there is a rule $X_1 \overset{p}{\hookrightarrow} \beta$ with $p > 0$ and $\beta(X_2) \geq 1$.

- If $k > 2$, the successor configuration is still unstable, but with probability at least p its corresponding path in G has length at most $k-1$.

– If $k = 2$, we have with probability at least p that the successor configuration α' satisfies $\alpha'(X_2) > \alpha(X_2) < K$.

Observe that if an increase $\alpha'(X_2) > \alpha(X_2) < K$ as described in the case $k = 2$ happens, then the X_2-component will remain above $\alpha(X_2)$ as long as the cautious strategy is followed, because the cautious strategy will not schedule X_2 as long as the X_2-component is at most K. Moreover, such increases can happen only finitely often before all types are saturated. It follows that with probability 1 a stable configuration will be reached eventually.

It is important to note that for cautious strategies the way how ties are resolved does not matter. Furthermore, for unstable $\alpha \in \mathbb{N}^\Gamma$ it does not matter if arbitrary types $X \in \Gamma$ with $\alpha(X) > K$ are scheduled in between. More precisely, for unstable α, consider two schedulers, say σ_1 and σ_2, that both follow a cautious strategy but may schedule other types X with $\alpha(X) > K$ in between. If for each type $X \in \Gamma$ the same probabilistic outcomes occur in the same order when following σ_1 and σ_2, respectively, then the resulting stable configurations β_1 and β_2 satisfy $\lfloor \beta_1 \rfloor = \lfloor \beta_2 \rfloor$. In other words, differences can only occur in saturated types.

Recall that in the finite-state MDP we need to redirect those transitions that lead to unstable configurations. We do that in the following way. For unstable α, let

$$T_\alpha := \{\lfloor \beta \rfloor \mid \beta \text{ is stable and reachable from } \alpha \text{ using a cautious strategy}\} \subseteq Q$$

and let $p_\alpha : T_\alpha \to (0, 1]$ be the corresponding probability distribution. As argued above, p_α does not depend on the particular choice of the cautious strategy. However, for the construction of the finite-state MDP one does not need to compute p_α, because for reachability with probability 1 in a finite-state MDP the exact values of nonzero probabilities do not matter. Note that we have $p_\alpha(q) > 0$ for all $q \in T_\alpha$. So if a scheduling action in the finite-state MDP would lead, in the infinite-state MDP, to an unstable configuration α with probability p_0, then in the finite-state MDP we replace this transition by transitions to T_α, each with probability $p_0/|T_\alpha|$. As argued above, this reflects a cautious strategy (which is optimal) of the scheduler in the original infinite-state MDP for states outside of Q.

This redirecting needs to be done for all unstable α that are reachable from Q within one step. There are only finitely many such α.

The overall decision procedure is thus as follows:

1. Construct the finite-state MDP with Q as set of states as described.
2. Check whether $\lfloor \alpha \rfloor$ is winning with probability 1, where the target set is $Q \cap F$.

If $\lfloor \alpha \rfloor$ is winning with probability 1, then there is a deterministic and memoryless scheduler. This scheduler can then be extended for the infinite-state MDP by a cautious strategy, resulting in a deterministic and memoryless scheduler.

It remains to show how T_α can be computed. We compute T_α using the decidability of the reachability problem for Petri nets. We construct a Petri net

from $\mathcal{S}$ that simulates cautious behaviour of the scheduler in unstable configurations. The set of places of the Petri net is $P := \Gamma \cup \{S_X \mid X \in \Gamma\}$, where the S_X are fresh symbols. The intention is that a configuration $\alpha \in \mathbb{N}^P$ with $\alpha(S_X) = 1$ indicates that X is saturated.

For the transitions of the Petri net we need some notation. For $\alpha, \beta \in \mathbb{N}^P$ we write $\alpha \hookrightarrow_\bullet \beta$ to denote a transition whose input multiset is α and whose output multiset is β. For $X \in P$ and $i \in \mathbb{N}$ we write X^i to denote $\alpha \in \mathbb{N}^P$ with $\alpha(X) = i$ and $\alpha(Y) = 0$ for $Y \neq X$. For $\alpha \in \mathbb{N}^\Gamma$ we also write α to denote $\alpha' \in \mathbb{N}^P$ with $\alpha'(X) = \alpha(X)$ for $X \in \Gamma$ and $\alpha'(X) = 0$ for $X \notin \Gamma$.

We include transitions as follows. For each $X \hookrightarrow \beta$ we include $X^{K+1} \hookrightarrow_\bullet X^K + \beta$. This makes sure that the transition $X \hookrightarrow \beta$ "inherited" from $\mathcal{S}$ is only used "cautiously", i.e., in the presence of more than K processes of type X. For each strongly connected component $\{X_1, \ldots, X_k\} \subseteq \Gamma$ of the graph G from the beginning of the proof, we include a transition $X_1^K + \cdots + X_k^K + \gamma \hookrightarrow_\bullet \{S_{X_1}, \ldots, S_{X_k}\} + \gamma$ with $\gamma = S_{Y_1} \ldots S_{Y_\ell}$ for $\{Y_1, \ldots, Y_\ell\} = \left(\bigcup_{i=1}^k succ(X_i)\right) \setminus \{X_1, \ldots, X_k\}$. This reflects the definition of "saturated": the types of a strongly connected component $\{X_1, \ldots, X_k\}$ are saturated in a configuration $\alpha \in \mathbb{N}^\Gamma$ if and only if $\alpha(X_i) \geq K$ holds for all $1 \leq i \leq k$ and all successors are saturated. We also include transitions that "suck out" superfluous processes from saturated types: $\{S_X, X\} \hookrightarrow_\bullet \{S_X\}$ for all $X \in \Gamma$.

For a configuration $q \in Q$ we define $\langle q \rangle \in \mathbb{N}^P$ as the multiset with $\langle q \rangle(X) = q(X)$ and $\langle q \rangle(S_X) = 0$ for $X \notin Sat(q)$, and $\langle q \rangle(X) = 0$ and $\langle q \rangle(S_X) = 1$ for $X \in Sat(q)$ (and hence $q(X) = K$). By the construction of the Petri net we have for all unstable $\alpha \in \mathbb{N}^\Gamma$ and all $q \in Q$ that a stable configuration $\beta \in \mathbb{N}^\Gamma$ with $\lfloor \beta \rfloor = q$ is reachable from α in $\mathcal{S}$ using a cautious strategy if and only if $\langle q \rangle$ is reachable from α in the Petri net. It follows that T_α can be computed for all unstable α. □

Theorem 11. *Given a pBPP $\mathcal{S} = (\Gamma, \hookrightarrow, Prob)$, an upward-closed set F, a number $k \in \mathbb{N}$, and a configuration $\alpha_0 \in \mathbb{N}^\Gamma$, it is decidable whether for all k-fair schedulers σ we have $\mathcal{P}_\sigma(\alpha_0 \models \Diamond F) = 1$.*

Proof. We extend the state space from $\mathbb{N}^\Gamma$ to $\mathbb{N}^\Gamma \times \mathbb{N}^\Gamma$. A configuration $(c, a) \in \mathbb{N}^\Gamma \times \mathbb{N}^\Gamma$ contains the multiset $c \in \mathbb{N}^\Gamma$ of current processes (as before), and an *age* vector $a \in \mathbb{N}^\Gamma$ indicating for each $X \in \Gamma$ how many steps ago X was last scheduled. We take $a(X) = 0$ for those X with $c(X) = 0$. An MDP $\mathcal{D}'(\mathcal{S})$ can be defined on this extended state space in the straightforward way: in particular, in each step in which an $X \in \Gamma$ with $c(X) > 0$ is not scheduled, the age $a(X)$ is increased by 1. We emphasize this extension of the state space does *not* result from changing the pBPP $\mathcal{S}$, but from changing the MDP induced by $\mathcal{S}$.

An age $a(X) \geq k$ indicates that X was not scheduled in the last k steps. So a run $(c_0, a_0)(c_1, a_1) \ldots$ is fair if and only if $a_i(X) < k$ holds for all $i \in \mathbb{N}$ and all $X \in \Gamma$. Define $G := (F \times \mathbb{N}^\Gamma) \cup (\mathbb{N}^\Gamma \times \{k, k+1, \ldots\}^\Gamma)$. There is a natural bijection between the k-fair runs in $\mathcal{D}(\mathcal{S})$ avoiding F and all runs in $\mathcal{D}'(\mathcal{S})$ avoiding G. So it suffices to decide whether there exists a scheduler σ for $\mathcal{D}'(\mathcal{S})$ with $\mathcal{P}_\sigma((\alpha_0, (0, \ldots, 0)) \models \Diamond G) < 1$.

To decide this we consider a turn-based game between two players, Scheduler ("he") and Probability ("she"). As expected, in configuration $\alpha \in \mathbb{N}^\Gamma \times \mathbb{N}^\Gamma$ player Scheduler selects a type $X \in \Gamma$ with $\alpha(X) \geq 1$ and player Probability picks β with $X \hookrightarrow \beta$, leading to a new configuration $T(\alpha, X \hookrightarrow \beta)$, where $T(\alpha, X \hookrightarrow \beta) \in \mathbb{N}^\Gamma \times \mathbb{N}^\Gamma$ denotes the configuration obtained from α by applying $X \hookrightarrow \beta$ according to the transitions of $\mathcal{D}'(\mathcal{S})$. Despite her name, player Probability is not bound to obey the probabilities in $\mathcal{S}$; rather she can pick β with $X \hookrightarrow \beta$ as she wants. The goal of Scheduler is to avoid G; the goal of Probability is to hit G.

We show that in this game one can compute the winning region for Probability. We define sets $W_0 \subseteq W_1 \subseteq \ldots$ with $W_i \subseteq \mathbb{N}^\Gamma \times \mathbb{N}^\Gamma$ for all $i \in \mathbb{N}$: define $W_0 := G$ and for all $i \in \mathbb{N}$ define

$$W_{i+1} := W_i \cup \{\alpha \in \mathbb{N}^\Gamma \times \mathbb{N}^\Gamma \mid \forall X \in \Gamma \; \exists \beta : X \hookrightarrow \beta \text{ and } T(\alpha, X \hookrightarrow \beta) \in W_i\}.$$

For all $i \in \mathbb{N}$ we have that W_i is the set of configurations where Probability can force a win in at most i steps. As $W_0 = G$ is upward-closed with respect to the componentwise ordering $\preceq$ on $\mathbb{N}^\Gamma \times \mathbb{N}^\Gamma$, all W_i are upward-closed with respect to $\preceq$. Considering the minimal elements of each W_i, it follows by Dickson's lemma that for some i we have $W_i = W_{i+1}$ and hence $W_j = W_i$ for all $j \geq i$. Then $W := W_i$ is the winning region for Probability. One can compute the minimal elements of W by computing the minimal elements for each $W_1, \ldots, W_i = W$. Define $\widetilde{G} := (\mathbb{N}^\Gamma \times \mathbb{N}^\Gamma) \setminus W$, the winning region for Scheduler, a downward-closed set.

Next we show that for $\alpha_0 \in \mathbb{N}^\Gamma \times \mathbb{N}^\Gamma$ we have that there exists a scheduler σ for $\mathcal{D}'(\mathcal{S})$ with $\mathcal{P}_\sigma(\alpha_0 \models \Diamond G) < 1$ if and only if in $\mathcal{D}'(\mathcal{S})$ there is a path from α_0 to $\widetilde{G}$ avoiding G. For the "if" direction, construct a scheduler σ that "attempts" this path. Since the path is finite, with positive probability, say p, it will be taken. Once in $\widetilde{G}$, the scheduler σ can behave according to Scheduler's winning strategy in the two-player game and thus avoid G indefinitely. Hence $\mathcal{P}_\sigma(\alpha_0 \models \Diamond G) \leq 1 - p$. For the "only if" direction, suppose that $\widetilde{G}$ cannot be reached before hitting G. Then regardless of the scheduler the play remains in the winning region W of Probability in the two-player game. Recall that $W = W_i$ for some $i \in \mathbb{N}$, hence regardless of the scheduler with probability at least $p_{min}^i > 0$ the set G will, at any time, be reached within the next i steps, where $p_{min} > 0$ is the least positive probability occurring in the rules of $\mathcal{S}$. It follows that for all schedulers σ we have $\mathcal{P}_\sigma(\alpha_0 \models \Diamond G) = 1$.

It remains to show that it is decidable for a given $\alpha_0 \in \mathbb{N}^\Gamma \times \mathbb{N}^\Gamma$ whether there is path to $\widetilde{G}$ avoiding G. But this can be done using a Karp-Miller-style algorithm as in the proof of Theorem 4. In particular, the following analogue of Lemma 5 holds:

Lemma 17. *Let $\alpha_0 \in (\mathbb{N}^\Gamma \times \mathbb{N}^\Gamma) \setminus G$ and let $\gamma \in \mathbb{N}^\Gamma \times \mathbb{N}^\Gamma$. Let $\alpha_0 \to \alpha_1 \to \ldots \to \alpha_k$ be a shortest path in $\mathcal{D}'(\mathcal{S})$ with $\alpha_0, \ldots, \alpha_k \notin G$ and $\alpha_k \preceq \gamma$. Then for all i, j with $0 \leq i < j \leq k$ we have $\alpha_i \not\preceq \alpha_j$.*

As argued in the proof of Theorem 4 one can build a (finite) tree of configurations reachable from α_0 via non-G-configurations and prune it whenever the path

$\alpha_0 \to \ldots \to \alpha_k$ from the root to the current node α_k contains a configuration α_j (where $0 \leq j < k$) with $\alpha_j \preceq \alpha_k$. If there is a path from α_0 to $\widetilde{G}$ via non-G-configurations, then the algorithms finds one. □

D Proofs of Section 6

Theorem 13. *Let $\mathcal{S} = (\Gamma, \hookrightarrow, Prob)$ be a pBPP. Let $F \subseteq \mathbb{N}^\Gamma$ be a semilinear set. Let $\alpha_0 \in \mathbb{N}^\Gamma$. The following problems are undecidable:*

(a) Does $\mathcal{P}(\alpha_0 \models \Diamond F) = 1$ hold?
(b) Does $\mathcal{P}_\sigma(\alpha_0 \models \Diamond F) = 1$ hold for all 7-fair schedulers σ?
(c) Does $\mathcal{P}_\sigma(\alpha_0 \models \Diamond F) = 1$ hold for some scheduler σ?

Proof.

(a) We will give a reduction from the complement of the control-state reachability problem for (deterministic) 2-counter machines (with counters X and Y). Given a 2-counter machine $M = (\mathcal{Q}, \Delta, q_0, q_F)$, we want to check if there is no computation from configuration $(q_0, 0, 0)$ to any configuration in $\{q_F\} \times \mathbb{N}^2$ in M. We will construct a pBPP $\mathcal{S} = (\Gamma, \hookrightarrow, Prob)$ that "simulates" M, a semilinear set $F \subseteq \mathbb{N}^\Gamma$, and a configuration $\alpha_0 \in \mathbb{N}^\Gamma$ such that $(q_0, 0, 0) \not\to^*_M \{q_F\} \times \mathbb{N}^2$ iff $\mathcal{P}(\alpha_0 \models \Diamond F) = 1$.
W.l.o.g., we assume that there is no transition from q_F in Δ. Define

$$\Gamma = \mathcal{Q} \cup \{q_{bad}\} \cup \{X_+, X_-, Y_+, Y_-\} \cup \{(\theta_1, \theta_2) \mid \exists q, q', c_1, c_2 : \langle (q, \theta_1, \theta_2), (q', c_1, c_2) \rangle \in \Delta\},$$

where $q_{bad}, X_+, X_-, Y_+, Y_- \notin \mathcal{Q}$ and $\theta_1 \in \{X = 0, X > 0\}$ (resp. $\theta_2 \in \{Y = 0, Y > 0\}$) is a zero test for first (resp. second) counters. Intuitively, in any pBPP configuration α, the number $\alpha(X_+) - \alpha(X_-)$ (resp. $\alpha(Y_+) - \alpha(Y_-)$) denotes the value of the first (resp. second) counter. For each $n \in \mathbb{Z}$, let $\mathrm{sign}(n) = +$ if $n \geq 0$; otherwise, let $\mathrm{sign}(n) = -$. For each transition rule $(q, \theta_1, \theta_2) \to (q', c_1, c_2)$ in Δ, we add to $\mathcal{S}$ a transition $q \hookrightarrow q'(\theta_1, \theta_2) X^{c_1}_{\mathrm{sign}(c_1)} Y^{c_2}_{\mathrm{sign}(c_2)}$, and a transition $(\theta_1, \theta_2) \hookrightarrow \varepsilon$. For each $Z \in (\mathcal{Q} \setminus \{q_F\}) \cup \{X_+, X_-, Y_+, Y_-\}$, we also add $Z \hookrightarrow q_{bad}$. Finally, we add $q_F \hookrightarrow q_F$. For each transition in $\hookrightarrow$, we do not actually care about its actual probability; we simply set it to be strictly positive.
We now define F to be the union of the following Presburger formulas:

1. $\neg \bigwedge_{(\theta_1,\theta_2)\in\Gamma} ((\theta_1, \theta_2) > 0 \to (\theta_1[(X_+ - X_-)/X] \wedge \theta_2[(Y_+ - Y_-)/Y]))$
2. $\neg \bigwedge_{(\theta_1,\theta_2)\in\Gamma} \Big((\theta_1, \theta_2) > 0 \to \Big((\theta_1, \theta_2) = 1 \wedge \bigwedge_{(\theta'_1,\theta'_2)\in\Gamma\setminus\{(\theta_1,\theta_2)\}} (\theta'_1, \theta'_2) = 0\Big)\Big)$
3. $q_{bad} > 0 \wedge q_F = 0$

where $\theta_i[Z/Z']$ denotes replacing every occurrence of variable Z' in θ_i with the term Z. The first conjunct above encodes bad configurations in $\mathcal{S}$ that are visited if the simulation of M by $\mathcal{S}$ is not faithful (i.e., counter tests are violated but the transitions are still executed). Since F is expressible in Presburger Arithmetic, it follows that it is a semilinear set.

We now prove the correctness of the reduction, i.e., that $(q_0,0,0) \not\rightarrow^*_M \{q_F\}\times\mathbb{N}^2$ iff $\mathcal{P}(\alpha_0 \models \Diamond F) = 1$. Suppose that $(q_0,0,0) \rightarrow^*_M (q_F,n_1,n_2)$ is witnessed by some path π, for some $n_1,n_2 \in \mathbb{N}$. This implies that, for some $c_1,c_1',c_2,c_2' \in \mathbb{N}$ such that $n_1 = c_1 - c_1'$ and $n_2 = c_2 - c_2'$, there is a finite path π' from q_0 to $\alpha = q_F X_+^{c_1} X_-^{c_1'} Y_+^{c_2} Y_-^{c_2'}$ that avoids F.
This path is a faithful simulation of π, which removes each (θ_1,θ_2) as soon as it is introduced. Since $\alpha \in \widetilde{F}$, it follows that $\mathcal{P}(Run(\pi')) > 0$ and so $\mathcal{P}(q_0 \models \Diamond F) < 1$. Conversely, assume that $\mathcal{P}(q_0 \models \Diamond F) < 1$. It is easy to see that, for the pBPP $\mathcal{S}$ that we defined above, both $\mathcal{M}_{type}(\mathcal{S})$ and $\mathcal{M}_{proc}(\mathcal{S})$ when restricted to states that are reachable from $\{q_0\}$ are globally coarse. Thus, there exists a finite path π from q_0 to $q_F X_+^{c_1} X_-^{c_1'} Y_+^{c_2} Y_-^{c_2'} q_{bad}^c$ that avoids F (since each configuration in $\widetilde{F}$ is of this form). In particular: (1) each time a symbol of the form (θ_1,θ_2) is introduced in π, it is immediately removed, and (2) rules of the form $Z \rightarrow q_{bad}$ (where $Z \in \{X_+,X_-,Y_+,Y_-\}$) are executed only after q_F is reached. Therefore, it follows that the path π is a faithful simulation of M and so corresponds to a path $\pi' : (q_0,0,0) \rightarrow^*_M (q_F,n_1,n_2)$ for some n_1 and n_2. This completes the proof of correctness of the reduction.

(b) The proof is a straightforward adaptation of the previous proof: for all $Z \in \{X_+,X_-,Y_+,Y_-\}$, we add the rule $Z \hookrightarrow Z$, and remove $Z \hookrightarrow q_{bad}$. That it suffices to restrict to 7-fair schedulers is because there are a total of 6 types for $\mathcal{S}$.

(c) The reduction is again from the acceptance problem of 2-counter machines. As before, we are given a 2-counter machine $M = (\mathcal{Q},\Delta,q_0,q_F)$ and we want to check if there is no computation from configuration $(q_0,0,0)$ to any configuration in $\{q_F\}\times\mathbb{N}^2$ in M. Without loss of generality, we may assume that $q_0 \neq q_F$ and that there is no transition in Δ of the form $\langle q,(\theta_1,\theta_2),(q,c_1,c_2)\rangle$ (i.e. stay in the same control state). We now define the pBPP $\mathcal{S}$. The set Γ of process types is defined as follows:

$$\Gamma = (\mathcal{Q}\cup\Delta)\times\{\bullet,\circ\}\cup \{X_+,X_-,Y_+,Y_-,Z,V\}$$

Let $\bar{\bullet} = \circ$ and $\bar{\circ} = \bullet$. Let $J = \{X_+,X_-,Y_+,Y_-,Z,V\}$. The initial configuration is $\alpha_0 = \{(q_0,\bullet)\}\cup\{(q,\circ) : q \in \mathcal{Q}\setminus\{q_0\}\}\cup(\Delta\times\{\circ\})$. The rules are as follows:

- for each $r \in \mathcal{Q}$ and $I \in \{\bullet,\circ\}$, we have $(r,I) \hookrightarrow (r,\bar{I})Z$ and $(r,I) \hookrightarrow (r,\bar{I})ZV$.
- for each $t \in \Delta$, we have $(t,\bullet) \hookrightarrow (t,\circ)Z$ and $(t,\bullet) \hookrightarrow (t,\circ)ZV$.
- for each $t = \langle q,(\theta_1,\theta_2),(q',c_1,c_2)\rangle$, we have $(t,\circ) \hookrightarrow (t,\bullet)X^{c_1}_{\mathrm{sign}(c_1)}Y^{c_2}_{\mathrm{sign}(c_2)}Z$ and $(t,\circ) \hookrightarrow (t,\bullet)X^{c_1}_{\mathrm{sign}(c_1)}Y^{c_2}_{\mathrm{sign}(c_2)}ZV$
- for each $W \in J$, we have $W \hookrightarrow W$.

The semilinear target set F is defined as a conjunction of the following Presburger formulas:

- $(q_F,\bullet) > 1$,

- $V < 2$
- If $V = 1$, then all of the following hold:
  - $X_+ - X_- \geq 0$ and $Y_+ - Y_- \geq 0$
  - for each $r \in \mathcal{Q} \cup \Delta$, $(q, \bullet) + (q, \circ) = 1$
  - If $Z \equiv 0 \pmod 4$, then (a) for precisely one $q \in \mathcal{Q}$ we have $(q, \bullet) = 1$ and $(q', \circ) = 1$ for each $q \neq q' \in \mathcal{Q}$, and (b) for each $t \in \Delta$, we have $(t, \circ) = 1$.
  - If $Z \equiv 1 \pmod 4$, then (a) for precisely one $q \in \mathcal{Q}$ we have $(q, \bullet) = 1$ and $(q', \circ) = 1$ for each $q \neq q' \in \mathcal{Q}$, (b) for precisely one $t = \langle q, (\theta_1, \theta_2), q', (c_1, c_2)\rangle \in \Delta$, we have $(t, \bullet) = 1$ and for each $t \neq t' \in \Delta$, we have $(t, \circ) = 1$, and (c) $\theta_1[(X_+ - X_- - c_1)/X]$ and $\theta_2[(Y_+ - Y_- - c_2)/Y]$ hold.
  - If $Z \equiv 2 \pmod 4$, then (a) for all $q \in \mathcal{Q}$ we have $(q, \circ) = 1$, (b) for precisely one $t = \langle q, (\theta_1, \theta_2), q', (c_1, c_2)\rangle \in \Delta$, we have $(t, \bullet) = 1$ and for each $t \neq t' \in \Delta$, we have $(t, \circ) = 1$, and (c) $\theta_1[(X_+ - X_- - c_1)/X]$ and $\theta_2[(Y_+ - Y_- - c_2)/Y]$ hold.
  - If $Z \equiv 3 \pmod 4$, then (a) for precisely one $q \in \mathcal{Q}$ we have $(q, \bullet) = 1$ and $(q', \circ) = 1$ for each $q \neq q' \in \mathcal{Q}$, (b) for precisely one $t = \langle q', (\theta_1, \theta_2), q, (c_1, c_2)\rangle \in \Delta$, we have $(t, \bullet) = 1$ and for each $t \neq t' \in \Delta$, we have $(t, \circ) = 1$, and (c) $\theta_1[(X_+ - X_- - c_1)/X]$ and $\theta_2[(Y_+ - Y_- - c_2)/Y]$ hold.

The target set F above forces the scheduler to do a faithful simulation of the input counter machine. For example, $(q_0, 0, 0) \to (q_3, 1, 0)$ via the transition rule $t = (q_0, (X = 0, Y = 0), q_3, (1, 0))$ is simulated by several steps as follows (we omit mention of $(r, \circ)$, when $(r, \circ) > 0$):

$$(q_0, \bullet) \hookrightarrow (q_0, \bullet)(t, \bullet)ZX_+ \hookrightarrow (t, \bullet)Z^2X_+ \hookrightarrow (q_3, \bullet)(t, \bullet)Z^3X_+ \hookrightarrow (q_3, \bullet)Z^4X_+$$

As soon as the scheduler deviates from faithful simulation, the Probability player can choose the rule that spawns V, which does not take us to F. For the next non-looping move (i.e. not of the form $W \hookrightarrow W$, which does not help the scheduler), Probability can spawn another V which takes us to $\widetilde{F}$, which prevents the scheduler from reaching F forever. Conversely, Probability cannot choose to spawn V if Scheduler performs a correct simulation; for, otherwise, F will be reached in one step. Therefore, this shows that there exists a scheduler σ such that $\mathcal{P}_\sigma(\alpha_0 \models \Diamond F) = 1$ iff the counter machine can reach q_F from $(q, 0, 0)$. □